Evaluating Behavioral and Neuroimaging Data on Past Tense Processing

Mark S. Seidenberg

James H. Hoeffner

Neuroscience Program and

Departments of Psychology and Linguistics

University of Southern California

Running head: Past tense processing

Address for correspondence:
Mark S. Seidenberg
Neuroscience Program
University of Southern California
Los Angeles, CA 90089-2520

email: marks@gizmo.usc.edu

fax: 213-740-5687

phone: 213-740-9174

Abstract*

Jaeger, Lockwood, Kemmerer, Van Valin, Murphy, and Khalak (1996) reported an experimental study that provided reaction time and PET neuroimaging data said to support Pinker’s (1991) theory of inflectional morphology in which rule-governed forms and exceptions are processed by separate mechanisms. The results were also taken as evidence against connectionist accounts in which a single processing system generates both types of forms. We provide a critical analysis of the study that yields three main conclusions: First, Jaeger et al.'s data do not provide strong evidence that rule-governed forms and exceptions are processed in separate brain regions. Second, there are problems with the design of the study that contaminate critical comparisons between conditions. The results therefore afford alternative interpretations related to experiment-specific factors rather than the regular-irregular distinction. Third, the dissociations between rule-governed forms and exceptions observed in studies such as Jaeger et al.'s can be accommodated by the connectionist theory. We conclude by offering suggestions for future research that would overcome the major limitations of this study and provide more decisive evidence bearing on the issues.

* Research supported by National Institute of Mental Health grants PO1-47566 and KO2-01188. The manuscript benefitted from the comments of the editor and three anonymous reviewers, whom we thank. We also thank Joe Allen (USC) and Jay McClelland (CMU) for helpful discussions. Jim Hoeffner is now at the department of psychology, University of Memphis.

1. Introduction. The formation of the past tense of verbs in English has been the focus of a long-running debate about the relevance of connectionist concepts to understanding language. 'Connectionism' refers to a body of ideas about knowledge representation, acquisition, and processing that have been applied to a broad range of issues concerning language and other aspects of cognition (Rumelhart & McClelland 1986). Connectionist theories attempt to explain phenomena in terms of networks of simple, neuron-like processing units whose behavior changes adaptively through learning. Such theories are typically instantiated as computational models that simulate detailed aspects of human performance (e.g., Seidenberg & McClelland 1989). Part of the attraction of the approach is that the same concepts are also being used in developing theories of the brain mechanisms underlying behavior (e.g.,Gluck & Rumelhart 1990; Elman et al. 1996), offering the possibility of a unified account of the two.

In one of the early studies within this approach, Rumelhart and McClelland (1986) implemented a model of some aspects of the acquisition of the English past tense. The model was a past tense generating device that took the present tense of a verb as input and produced the past tense as output. Its main interest lay in the fact that it could generalize to novel forms (e.g., wug-wugged). The model was therefore said to provide an alternative to approaches in which generativity results from the use of rules. Several limitations of this model were subsequently noted by Pinker and Prince (1988). These limitations were said to vitiate Rumelhart and McClelland's claim to have provided an alternative account of linguistic knowledge, and were used to call into question the relevance of connectionist concepts to understanding language: Pinker and his colleagues subsequently acquired a large amount of evidence from a variety of domains said to support a 'modified traditional theory' of the past tense employing two mechanisms: a rule used to generate regular forms (e.g., walk-walked and perform generalization (e.g., wug-wugged), and an 'associative net' used in generating irregular forms (e.g., take-took) (see Pinker 1991 for a summary). Connectionists have also presented modified models that addressed many of the limitations of Rumelhart and McClelland's original work (see, for example, Daugherty & Seidenberg 1992; MacWhinney & Leinbach 1991) and have challenged findings taken as supporting the dual-mechanism theory (e.g., Hoeffner, 1997).

Because the past tense of English is a minor aspect of language and because this controversy has continued without resolution for almost 10 years, it is worth remembering why these issues are important. Studies of the past tense (and inflectional morphology more generally) have historically played important roles in several core areas of linguistic research, including phonological theory (e.g., Halle & Mohanan 1985), language acquisition (e.g., Berko 1958; Bybee & Slobin 1982), adult psycholinguistics (e.g, Bybee & Moder 1983), diachronic change (e.g., Hare & Elman 1995), and the genetic (Rice 1996) and neurobiological (Ullman et al. 1997) bases of language. These studies illustrate the value of exploring a specific aspect of language in considerable detail. However, it is also true that if the connectionist approach were proved to be incapable of explaining something as simple as the past tense, this result would impose severe limitations on its relevance to other, more complex aspects of language.

The article by Jaeger, Lockwood, Kemmerer, Van Valin, Murphy and Khalak (1997; hereafter J et al.) seems to provide compelling evidence bearing on this controversy. Pinker's theory draws a strong distinction between rule-governed forms and exceptions: they are produced by separate subsystems that involve different types of knowledge representations (a rule, an 'associative net'), learning procedures (rule induction, rote learning), and processing mechanisms (rule application, lexical look-up). In the connectionist approach, a single computational system is used to generate all forms. The same principles of knowledge representation, learning, and processing therefore govern both types of forms. The essential idea is that a network trained to generate past tenses can find a set of weights that allow it to produce all forms. The network therefore does not embody a categorical distinction between 'rule-governed' and 'exception.' Using a single mechanism to generate all forms is intended to capture the fact that rule-governed forms and exceptions are not categorically different; they share structure in at least three respects. First, there are systematic correspondences between the present and past tense forms of both regular and irregular verbs. For example, in both <u>bake-baked</u> and <u>take-took</u>, the past tense retains the onset and coda of the present tense. The degree of overlap varies across verbs but there is some for all past tense alternations except the suppletive ones. Second, there are similarities between regular and irregular past tense forms. For example, crept and slept, which are nominally irregular, end in /t/, which is the regular past tense inflection in cropped, stepped and many other regular forms. Finally, there are subregularities among pools of irregular verbs, such as sing-sang/ring-rang, and blow-blew/throw-threw/grow-grew. These properties motivated earlier treatments such as Halle and Mohanan (1985) and Bybee (1985). The neural network approach shows how these partial regularities can be picked up in the course of training and encoded by the weights along with the regularities characteristic of the past tense rule. The weights that are derived from exposure to examples can also be used to generate past tenses for novel forms such as wug.[1]

J et al. presented the results of a positron emission tomography (PET) neuroimaging study providing evidence about brain activity during past tense generation. The data were interpreted as indicating that rule-governed and irregular forms activate different brain regions, consistent with Pinker's dual-mechanism theory and contradicting the connectionist approach. Data concerning subjects' latencies to produce past tenses provided additional support for this conclusion.

The main point of this article is that there are conceptual, methodological, and interpretive flaws in the J et al. study that undermine their claims. Although we agree that neuroimaging studies have the potential to provide important evidence bearing on cognitive and linguistic theories, this particular study does not succeed in this regard. We will discuss three main points: First, J et al.'s data do not provide strong evidence that rule-governed forms and exceptions are processed in separate brain regions. Second, the design of the study introduced confounding factors that contaminate critical comparisons between conditions. These affect both the imaging and latency data. The results therefore afford alternative interpretations related to experiment-specific factors rather than the regular-irregular distinction. Third, the assumption that dissociations between regular and irregular forms both implicate separate mechanisms and contradict the connectionist theory needs to be examined closely. The particular dissociations that

J et al. were seeking are compatible with both types of theories and so do not discriminate between them. The same issue arises in connection with other recent reports of such dissociations (Marslen-Wilson & Tyler 1997; Ullman et al 1997). In closing we suggest directions for future research, describing alternative hypotheses that would distinguish between the theories and methodologies appropriate for addressing them.

2. The PET Data. J et al.'s study provides data concerning brain activity while subjects were generating the past tenses for known, rule-governed forms (e.g., pull-pulled), irregular forms (sweep-swept) and nonce forms (baff-baffed).[2] Stimuli were presented in separate blocks of 46 items of each type, yielding the Regular, Irregular, and Nonce Past conditions, respectively. The PET images reflected average activity during a block of 46 trials of a single type. In two control conditions, subjects simply read words aloud rather than generating past tenses. The Read Verb control condition consisted of a mixture of regular and irregular verbs; in the Read Nonce control condition, subjects read aloud a different set of nonce verbs. Brain areas associated with generating past tenses for regular verbs were determined by subtracting the pattern in the Read Verb condition from the pattern in the Regular Past condition; analogous subtractions were performed for the Irregular and Nonce conditions. J et al. concluded that: 'there are major differences in both the location and the amount of brain activation in the regular vs. irregular tasks; specifically, overlapping yet partially distinct sets of brain areas are recruited for the computation of regular vs. irregular forms, with a much larger area of activation involved in the irregular computation. Thus...our data are consistent with the dual-systems theories and are problematic for the single-systems accounts.' (p. 484)

Table 1 presents the data in J et al.'s Table 4 (p. 467) but regrouped to make it easier to see the results for each verb condition. Three aspects of these data (highlighted in bold) were taken as providing crucial support for the dual-mechanism theory. First, part of the left dorsolateral prefrontal cortex was active in the Regular Past and Nonce Past conditions, but not the Irregular Past condition (after the relevant control conditions were subtracted out). This area was thought to be involved in generation by rule. Second, the left mid-temporal gyrus was active in the Irregular Past condition but not in the Regular Past or Nonce Past conditions (against subtracting out the relevant controls). This area was said to reflect retrieval of stored irregular forms. Finally, the left lateral orbito-frontal cortex was active in both the Irregular Past and Nonce Past conditions. J et al. argue that this area is involved in 'inhibiting default, high frequency responses, and is inactive during overlearned, practiced behaviors' (p. 483), the latter including generating a known rule-governed form.

Insert Table 1 About Here

This interpretation of the data raises two questions: First, what pattern of data would count as evidence against the dual-mechanism theory, and second, what is the basis for the attributions about what is being computed in different brain regions? The dual-mechanism theory suggests that regular and nonce verbs should pattern together and differ from irregular verbs. The first two findings mentioned above are consistent with this prediction: there were regions activated for rule-governed forms but not exceptions and vice versa. However, by the same logic we would not expect to see areas that are activated for both regular and irregular verbs or for irregular and nonce verbs, yet the data also yielded these patterns. In fact, there were unique brain regions activated for all three pairwise combinations of conditions, and there were areas activated by all three types of stimuli as well. Given the absence of any independent theory about why particular regions might be involved in processing particular types of stimuli, this data pattern cannot be taken as providing strong evidence for the dual-mechanism theory. The fact that some areas were activated for both irregulars and nonce items or for regulars and irregulars could just as well be taken as decisive evidence <u>against</u> the theory. In fact there are more areas activated in common by the irregulars and nonce words, which are supposed to be processed by separate mechanisms, than the regulars and nonce words, which are supposed to be generated by the rule.

J et al. discuss their data fully and attempt to provide interpretations of the activity in each brain region consistent with the dual-mechanism theory. However, their account turns on attributions about the functions associated with different brain regions that are highly questionable; they lack independent motivation and they are not internally consistent. For example, in explaining the activity in the left dorsolateral prefrontal cortex associated with the regular and nonce conditions, J et al. argue that this area of the brain is responsible for ‘using knowledge stored in long-term or short-term memory to formulate novel conceptual and behavioral sets in a willful, internally generated manner, as opposed to a purely stimulus-driven manner.’ (p. 482). They also state that this area of the brain has been shown to be involved in the ‘on-line formulation of intentional novel behaviors.’ (p. 488). This account suggests that applying linguistic rules is ‘willful’ rather than merely ‘stimulus-driven’ but this characterization of the use of linguistic knowledge conflicts with the views of most psycholinguists, who emphasize the rapid, automatic, unconscious nature of the processes involved in comprehension and production. Moreover, whereas generating past tenses for nonsense words may be ‘novel’ behavior, generating past tenses for known regular verbs is not. To the contrary, it is a highly practiced behavior. Why a brain area would be activated for behaviors that are willful and novel (generating past tenses for nonce words) and familiar and unconscious (generating known past tenses for familiar words) is unclear and unsupported by any independent neurophysiological evidence. Below we offer an alternative account of the activity of this region that is motivated by independent evidence and helps to explain other aspects of J et al.'s data.

J et al.'s interpretation of the activity in the left dorsolateral prefrontal cortex creates other problems when it comes to explaining the finding that the irregular and nonce conditions activated a common region (the left superior frontal gyrus). J et al. claim that this area is involved in ‘inhibiting default, high-frequency responses, and is inactive during overlearned, practiced behaviors.’ (p. 482). This interpretation assumes that producing the regular inflection is a ‘default, high-frequency behavior’ that is an ‘overlearned, automatic response’ (p. 487) and must be actively suppressed in producing the irregular. However, the interpretation of the activity in the left dorsolateral prefrontal cortex discussed above assumed that generating the regular suffix was ‘intentional’ and ‘willful’ behavior. Moreover, although it is clear why the regular past tense mechanism might have to be suppressed in generating an irregular past tense, this reasoning does not extend to the nonce task, where the regular suffixation process needs to apply, not be suppressed. J et al. address this apparent contradiction by arguing that in the nonce task this area is not inhibiting the regular suffixation process per se but is "involved in inhibiting any number of possible irrelevant responses to this novel list of forms." (p. 487). The problem with this argument, aside from its post hoc nature, is that this brain area is now assumed to serve different functions in the processing of irregulars and nonce words. Thus, when the regular and nonce verbs activated a common brain region, it was taken as evidence for a single mechanism underlying the two but when the irregulars and nonce items activated a common area it was interpreted in terms of two mechanisms. This inconsistency illustrates the weak linkage between theory and data in J et al.'s account.

J et al.'s assumptions about localization of function should also be considered in light of two recent studies of brain-injured patients' impairments in verb morphology (Ullman et al. 1997; Marslen-Wilson & Tyler 1997). All three studies claimed to find support for the dual mechanism theory, but they yielded conflicting evidence with respect to where the putative subsystems are localized. Data from several types of brain injured patients led Ullman et al. to conclude that a frontal lobe/basal ganglia system is responsible for regular inflection, with irregulars processed by a separate system located in temporal-parietal cortex. However, J et al.'s data implicate various parts of the frontal lobes in both regular and irregular inflection. In the Generate Past-Read Verb comparisons, four areas in the left hemisphere's frontal lobes were active for regular verbs, but three partially overlapping areas of the frontal lobes were also active for the irregular verbs. Similarly, parts of the temporal and parietal lobes were active for irregulars, as Ullman et al. predict, but also regular and nonce verbs. Finally, J et al. did not observe above threshold activation of the basal ganglia in any condition. In Marslen-Wilson and Tyler's study, it was damage to the right hemisphere that played an important role in the impairment of the irregular system but neither Ullman et al nor J et al. ascribe a critical role to this region. In addition, the patients with impaired performance on regular items in Marslen-Wilson & Tyler's study had damage to left hemisphere temporal and parietal lobes, which according to Ullman et al. should have led to impaired irregular performance. Our view is that these inconsistencies are due in part to the fact that the regular-irregular distinction does not capture the relevant generalization about performance, which is affected by the frequency and consistency of the mapping between present and past tenses. In any case, the complex and partially conflicting patterns of results in these studies suggest that considerable additional research is needed in order to understand the brain bases of past tense representation and processing.

3. Methodological issues. Even if J et al.'s data had yielded clearly different patterns for rule-governed forms and exceptions, they could not be taken as decisive because of a serious flaw in the design of the study. The flaw relates to the presentation of different types of stimuli in homogeneous blocks. Each of the critical conditions (Regular Past, Irregular Past, Nonce Past) consisted of 46 items of that type. This blocked design is imposed by the low temporal resolution of the PET methodology, which requires subjects perform the same type of task over the interval during which the image is computed. The presentation of stimuli in homogeneous blocks introduces confounding factors that invalidate the comparisons between conditions.

In the Regular Past condition, subjects had to add the regular suffix 46 consecutive times. All that was required was determining which of the three allomorphs of the regular suffix (illustrated by jump-jumped, love-loved, pat-patted) was correct for each verb. These allomorphs are phonetically similar, differing only in the presence or absence of voicing and/or the insertion of an unstressed vowel. Choice of allomorph is determined by the final phoneme of the present tense verb. The 46 items in the Irregular Past condition differed in several respects. They form the past tenses in a wide variety of ways including several vowel changes (e.g., fall-fell; sit-sat; bite-bit), changing final /d/ to /t/ (e.g., build-built; send-sent), and both vowel and consonant changes (bring-brought; keep-kept; stand-stood; sell-sold). They overlap in more complex ways than the regulars: J et al.'s stimuli included verbs that rhyme and form their past tenses similarly (e.g., tell-told; sell-sold) but also verbs that rhyme and form their past tenses differently (e.g., bring-brought; sing-sang). There were also partially-overlapping trials such as feed-fed vs. sweep-swept (same vowel change, different treatment of final consonant) and some irregular present tenses rhyme with the past tenses of other irregulars in the list (e.g., hold rhymes with told and sold). Finally, generating the correct irregular past tense requires recognizing the present tense verb as a particular lexical item, not just its final phoneme.

Given these differences between regular and irregular verbs, presenting the stimuli in blocks introduced confounding factors that are certain to have affected performance. One confound is that the blocks created different kinds of priming effects. A large body of psychological research has addressed how performance on one trial in an experiment is affected by performance on similar or dissimilar preceding trials. Such priming effects are ubiquitous; they occur for adjacent items (e.g., Forster & Davis 1984) and for items separated by many intervening trials (e.g., Scarborough et al. 1977). Stimuli that are similar create facilitation; stimuli that are dissimilar create interference. To take a simple example, Taraban and McClelland (1987) showed that pronunciation of the word gave is facilitated if it is preceded by the rhyming word save but inhibited by have. Unless they are themselves the focus of study (as in Gordon, 1983), experimenters typically take pains to avoid repeating similar or identical stimuli within an experiment. J et al. could not do this, however, because of limitations of the PET methodology. Presenting stimuli in homogeneous blocks creates massive priming effects that differ across conditions. Whereas the high degree of similarity across trials in the Regular Past and Nonce conditions facilitated performance (as indicated by the latency data discussed below), the variability across trials in the Irregular Past creates interference. Readers can verify this by conducting the experiment themselves. Table 2 contains the first 23 irregular stems used by J et al. in the same order as they were presented to subjects. Read through the list generating the past tenses quickly and accurately. There should be a sense that forming the past tense differently on each trial creates interference. Table 2 also presents a list of regular verbs for comparison. Table 3 presents a list of irregular verbs grouped into families that pattern similarly. This grouping simplifies the task considerably. Note that J et al.'s Regular Past condition created even stronger grouping effects because all 46 items conformed to the single regular pattern.[3] Thus, the blocks of regular and irregular stimuli created different priming effects due to differences in the amount and degree of overlap across trials that are confounded with type of verb. Studies of priming effects using PET and functional magnetic resonance imaging (fMRI) have shown that different kinds of priming produce different patterns of brain activity (e.g., Gabrieli et al, 1996; Demb et al., 1995). Given J et al.'s design, it cannot be determined whether their PET images reflect different types of priming or verbs.

Insert Tables 2 and 3 About Here

A second consequence of presenting the stimuli in blocks is that the Regular and Nonce Past conditions differed from the Irregular Past condition in terms of how deeply stimuli had to be processed in order to generate correct responses. In the Regular and Nonce conditions, subjects knew in advance of each stimulus that the regular past tense suffix would be required. The form of the past tense was determined by phonological properties of the present tense verb. In the Irregular Past condition, however, subjects did not know in advance how the past tense would be formed; producing correct responses required identifying each present tense verb as a particular lexical item. The Regular and Nonce Past conditions therefore differed from the Irregular Past condition in terms of the depth of processing (Craik & Lockhart 1972) that was required: the Irregular condition required a deeper analysis of each item than the other conditions. This confound is important because other neuroimaging studies indicate that stimuli that are processed more deeply produce different brain activity than stimuli that are processed more superficially (e.g., Gabrieli et al. 1996). Note that in a design in which regular and irregular forms are randomly intermixed this confound is eliminated: if the subject does not know in advance whether the next stimulus will be regular or irregular verb, all verbs have to be processed equally deeply.

J et al. acknowledged that their subjects quickly realized that all the verbs in Regular Past condition were regular (p. 476) and expressed appropriate concern that this aspect of the study might have affected the results, noting that 'The obvious test of the influence of task format is to gather RT's to stimuli where an equal number of frequency-matched regular and irregular stems are randomly mixed.' (p. 476). As they observe, PET methodology is incompatible with such a design because it requires subjects to be exposed to many consecutive examples of a given type. Unfortunately, all this indicates is that PET is an inappropriate method for addressing the issues at hand. Acknowledging that there are confounds between the conditions that are likely to influence the results does not eliminate the problem. A better alternative would have been to consider using methods that do not impose the same restrictions as PET, one of which we discuss below.[4]

4. Anomalies in the latency data. Our assertions about the effects of blocking the stimuli are supported by J et al.'s latency data and help to explain some anomalous aspects of them. J et al.'s data can be compared to that from a similar study conducted by Seidenberg and Bruck (1990, hereafter SB; Seidenberg 1992).[5] In SB's study (see Table 4), subjects were presented with 40 regular and 40 irregular present tense verbs (equated in frequency) and asked to generate the corresponding past tense forms. The design was the same as J et al's except that the regular and irregular verbs were intermixed rather than presented in separate blocks. Comparisons between the two studies therefore provide direct evidence about the effects of blocked vs. mixed presentation.

Insert Table 4 About Here

Several aspects of the data indicate that presenting the stimuli in blocks had a huge impact on the results. J et al.'s subjects were much faster than SB's and the speed-up was largely due to the regular verbs, which yielded mean latencies of 524 ms and 925 msec in the J et al. and SB studies, respectively. This pattern is what would be expected if presenting stimuli in blocks allowed subjects to take advantage of prior knowledge that all the stimuli were going to be regularly inflected. The subjects in SB's study had net generation latencies (latency to generate past tense - latency to read verb aloud) of 350 ms for regular verbs and 455 ms for irregular verbs. In contrast, the mean generation latencies in these conditions in J et al.'s study were 38 ms and 238 ms. Thus, regular past tenses were produced approximately 10 times faster in the J et al. experiment. There was also a smaller speedup for the irregular verbs. These data strongly indicate that the subjects in the J et al. study were using response strategies conditioned on the presentation of stimuli in blocks. The fact that it only took 38 ms on average to generate regular past tenses is itself striking: other research suggests that it takes at least 100 msec to merely identify a word and fixation durations for words in texts typically average 200-250 msec (Rayner & Pollatsek 1994). It is hard to imagine that applying a morphological rule would take an order of magnitude less time to perform than these tasks unless subjects could exploit the knowledge that all items in the block were regular.

The data in J et al.'s nonce conditions are also inconsistent with the results of other studies. In most studies, generating responses for familiar items is much easier than generalizing to novel forms. For example, subjects are faster at generating the pronunciation of a familiar word such as must than a nonword like nust. In the J et al. study, however, subjects produced past tenses for nonsense words slightly faster than for familiar irregular words. Table 4 presents the data from the J et al. study and a similar experiment by Seidenberg and Bruck (1990).[6] Stimuli were again blocked in the J et al. study; in the SB study, they were intermixed with an equal number of known regular and irregular verbs. The SB data are means from 52 nonce items and 30 subjects. Subjects in the two studies took roughly the same amount of time to name the nonwords, but SB's subjects took far longer to generate their past tenses. Latencies in the SB study were longer because subjects did not know that the stimulus on every trial would end with the regular past tense. Subjects were simply told that the stimuli would include both familiar and unfamiliar forms and were instructed to generate past tenses for the unfamiliar ones 'as if' they were real verbs in the language.

The comparisons between these studies clearly point to large effects attributable to demand characteristics of the J et al. experiment. Presenting the stimuli in homogeneous blocks provides information that subjects can use in planning their responses. When the stimuli are all regular, subjects can respond knowing that all that is required is adding the correct phonologicallyconditioned regular suffix. The task is highly predictable and repetitious, yielding latencies that are little longer than those for simply saying the verbs aloud. Because the identity of the stimulus is not critical to producing the response, it matters little whether it is a familiar word or a nonword. Blocking the irregulars is also helpful insofar as subjects know that the regular past tense will not be required. However, the benefit is smaller than in the regular case because subjects must still process each verb deeply enough to determine its past tense; moreover, forming the past tense differently on each trial produces interference that slows performance. These effects are experiment-specific and have little relevance to normal performance, in which speakers and listeners routinely switch between regular and irregular pasts and other tenses.

To summarize, effects due to repetitions of stimuli and other nonrandom aspects of stimulus presentation are robust, well-studied phenomena in experimental psychology. They are so prominent that in studies employing standard reaction time measures researchers use experimental designs that avoid them or allow their effects to be assessed. There are literally dozens of experiments in the literature examining the processing of different types of words (e.g., abstract-concrete; homophone-nonhomophone; high-low frequency; familiar-unfamiliar; regular-irregular pronunciation; word-nonword) and the stimuli are never blocked by type unless the explicit purpose of the experiment is to examine blocking itself. The reasons why researchers avoid this design feature are illustrated by J et al.'s study: repeating stimuli of a given type creates different kinds of priming effects across conditions and encourages condition-specific strategies.

In presenting stimuli in homogeneous blocks, J et al. were following the standard PET methodology pioneered by Raichle and his colleagues (e.g., Petersen, Fox, Posner, Mintun, & Raichle, 1988). (199x). Every methodology has its limitations; in the case of PET, it is the slow temporal resolution of the measure, which necessitates averaging across a block of similar trials. The use of PET to localize cognitive functions has generated considerable controversy about basic aspects of the methodology (Roskies 1994), including what brain events PET actually measures (Barinaga, 1997), the validity of the subtractive methodology (Sergent 1994), and the reliability of localization claims (Poeppel 1996). As with other methodologies, the utility of the approach depends on the kinds of questions being addressed. Raichle and his colleagues' classic research, for example, involved comparisons between conditions such as passively looking at actual words vs. nonsense words or reading verbs vs. generating an associated action. In these cases, each trial carries very little information about what will occur on other trials and subjects cannot formulate response strategies that work for one condition but not another. Hence the ceteris paribus condition that is the foundation of experimental research was largely met and the studies yielded coherent findings. In J et al.'s study, however, presenting the stimuli in homogeneous blocks introduced several confounding factors that invalidate the attribution of the observed effects to the nominal independent variable, type of verb.

5. An alternative interpretation: more difficult = more activation. Once the consequences of presenting stimuli in blocks are recognized, it is apparent that J et al.'s data afford a completely different interpretation that is consistent with the single mechanism theory. Because of the way the stimuli were selected and presented, the regular, irregular and nonce conditions differed in difficulty. The regular verbs were easiest, the nonce verbs were slightly harder, and the irregular verbs were much harder. A simple interpretation is that the observed patterns of brain activity reflected the relative difficulty of these conditions. The data are therefore consistent with the single-mechanism theory in which performance is determined by the degree of difficulty in the mapping between present and past tense, not type of verb.

The fact that there is a confound between type of verb and level of difficulty is crucial to understanding the PET data and claims about localization of function. J et al.'s data indicate that the irregulars produced the longest generation latencies and the most errors, the regulars yielded the fastest responses and fewest errors and the nonce verbs fell in between. These differences in difficulty are not intrinsic to these types of stimuli; it would be possible to equate the conditions in terms of difficulty by using different items (e.g., harder regular verbs or easier irregulars) and that is necessary to allow valid comparisons between conditions. Research by Just et al. (1996) has indicated that task difficulty is correlated with amount of brain activity. Just et al. conducted an fMRI study of language comprehension that involved three main conditions in which subjects read sentences varying in syntactic complexity. The relative difficulty of the stimuli was independently established from reaction time data. The main finding was that the difficulty of the materials was correlated with amount of brain activity: the harder stimuli produced activation over a broader area than the easier stimuli.

J et al.'s study replicates this result: there was a high correlation (.835) between the difficulty of the conditions (as indicated by response latencies) and the amount of brain activity they elicited. The irregular condition activated the most areas (15, 3 of which reached significance by the strictest thresholding method) and a total of 2,912 pixels were active after the baseline condition was subtracted. The nonce condition produced activation across 14 areas (none of which reached significance by the strictest criterion) and 1,368 pixels. The regular condition produced the least activation, only 728 active pixels in 7 areas, none of which reached significance by either the strictest or second-strictest criteria.

Consider now the three findings that J et al emphasize in their discussion of the Generate Past-Read Verb subtractions. One area, the left mid-temporal gyrus (area 21), was active for irregulars but not for nonce or regulars. J et al.'s interpretation is that this area was responsible for irregular inflection. However, this activity may simply reflect the fact that the irregular condition was the most difficult, and more difficult conditions produce more brain activity over a broader area. The inference that the area was involved in processing irregular verbs would only be valid if the conditions had been equated in terms of difficulty.

The second area J et al. emphasized is the left superior frontal gyrus (area 10), which was active for the irregulars and nonce verbs but not for the regulars, leading J et al to speculate that it is responsible for inhibiting the regular, 'default' process. Again, however, the irregular and nonce conditions were harder than the regular condition and harder tasks activate more areas. Moreover, the dual-mechanism theory assumes that irregulars and nonce items are processed by separate mechanisms; any rationale for why they activate the same brain region cannot provide direct evidence for the separate mechanisms, unless it is assumed that separate mechanisms can occupy the same region of the brain, which is contrary to the logic of using PET to identify distinct functions with distinct loci.

The third area J et al. emphasized is the left dorsolateral prefrontal cortex (area 46), which was active for regular and nonce verbs but not for the irregulars. This finding cannot be due to the relative difficulty of conditions because the regular and nonce conditions were easier than the irregular condition. There is, however, a compelling alternative interpretation of this result for which there is independent neurophysiological evidence. Recall the earlier suggestion that presenting the rule-governed stimuli in homogeneous blocks enabled a simple experiment-specific response strategy because subjects knew in advance that the regular suffix was required. This strategy requires maintaining this information in working memory. Classic studies with monkeys (Goldman-Rakic 1988) implicate the left dorsolateral prefrontal cortex in working memory. The area is active when, for example, a monkey needs to remember a particular spatial location over a delay period. We suggest that the left dorsolateral prefrontal cortex was playing a similar role in J et al.'s study. Neurons in this area of the monkey brain are thought to keep active a trace of the spatial location that the monkey must remember in order to be rewarded. Similarly, in J et al.'s Regular Past and Nonce Past tasks this area of the brain was likely to be involved in remembering that the regular suffix was required.

In summary, the presentation of stimuli in blocks in the J et al. study created differences in degree of difficulty across conditions that contributed to the patterns they obtained. These observations suggest that similar results would obtain for conditions consisting of verbs of a single type (e.g., regular) that exhibited analogous differences in difficulty (because, for example, they varied in frequency).

6. Do dissociations implicate separate subsystems? To this point the question we have addressed is whether J et al.'s data demonstrate differences in how rule-governed verbs and exceptions are processed. We now need to consider whether such differences, should they exist, necessarily favor the dual-mechanism theory. In interpreting their data, J et al. followed the lead

of Pinker, who has said 'Because it categorically distinguishes regular from irregular forms, the rule-association hybrid predicts that the two processes should be dissociable from virtually every point of view.' (Pinker 1991:532) The assumption that only the dual-mechanism theory predicts such differences has led to a vigorous hunt for them in a broad variety of domains. The reasoning underlying the interpretation of such dissociation data needs to be examined closely, however. It is true that the differences between two types of items can be explained in terms of two mechanisms dedicated to processing them, simply by investing each mechanism with whatever characteristics the data demand. However, the converse--that such differences are incompatible with the connectionist approach--is not true. There is an intuition that because connectionist models use a single mechanism to generate all forms, they provide no basis for predicting any differences between the forms that are classified as 'regular' and 'irregular' by the other theory. However, in the networks, verb alternations differ in terms of how hard they are to acquire and process and how vulnerable they are to different types of damage. These aspects of model performance provide a basis for explaining the kinds of dissociations that have been taken as favoring the single-mechanism account. Extensive research on this issue has led to one of the most important findings to have emerged from the first generation of connectionist models of disordered language and cognition: that dissociation data thought to transparently implicate separate mechanisms actually do not (Plaut et al., 1996; Plaut 1997).

The evidence bearing on these issues derives from studies of both the past tense and closely related phenomena concerning the pronunciation of visually-presented words. Studies of verb morphology such as J et al.'s involve presenting a verb and generating its past tense. Studies of word reading involve presenting a word and generating its pronunciation. Both bodies of knowledge are what we have termed 'quasiregular' systems in which most patterns can be described by rules but there are many exceptions that deviate from the rules in varying degrees (Seidenberg & McClelland 1989; Seidenberg 1992). In past tense morphology, there are rule-governed forms such as <u>bake-baked</u> and exceptions such as <u>take-took</u>. In word pronunciation, there are rule-governed forms such <u>gave</u> and <u>mint</u> and exceptions such as <u>have</u> and <u>pint</u>. In both domains the exceptions tend to be higher frequency items and form overlapping clusters. In both domains generalization is assessed by using nonce forms such as wug.

There are close similarities between the competing theories in the two domains as well. Pinker's dual-mechanism account of the past tense is formally identical to Coltheart et al.'s (1993) 'dual-route' theory of pronunciation. Both consist of a rule component and an associative network used to generate the exceptions. In both domains there is a competing connectionist theory in which all forms are generated by a single network. There is a vast amount of research on normal and disordered word pronunciation and both connectionist and dual-mechanism theories have been worked out in more detail than in the case of the past tense. Given the close similarities between the theories and data in the two domains, the word pronunciation research can provide useful insight about the past tense.

Consider again the J et al. study. Leaving aside the methodological issues raised above, the study yielded different patterns for regular and irregular verbs. Can such data be taken as evidence for two mechanisms? The answer is negative because this 'dissociation' is reliably exhibited by connectionist models that lack separate subsystems dedicated to different types of verbs. J et al.'s stimuli consisted of lists of regular and irregular verbs that were roughly equated in terms of frequency of occurrence in the language. The verbs in both groups ranged in frequency from very common (expected to occur about once every 1,000 words) to much less common (expected to occur about once every 100,000 words). We know from behavioral studies and connectionist models in both the word pronunciation and past tense domains that, other factors being equal, these kinds of stimuli yield an overall advantage for the regulars (Daugherty & Seidenberg 1992; Plaut et al. 1996). In brief, the weights in the network reflect the aggregate effects of exposure to all of the patterns in the corpus (e.g., present-past tense pairs). Performance on any given pattern depends on how it overlaps with other patterns. 'Rule-governed' present-past alternations benefit from the fact that there is a large body of other items that behave similarly ('friends') and few if any that behave dissimilarly ('enemies'). 'Irregular' alternations have fewer friends and many enemies. The weights in the network come to reflect these facts about the ensemble of present-past tense mappings. Thus, the overall differences that J et al. observed do not demand separate mechanisms. Moreover, as argued above, the fact that the irregulars were harder is relevant to interpreting the patterns of brain activity that were observed.

Our view is that this situation is typical: findings that have been claimed to favor the dual-mechanism theory can be derived (sometimes in nonobvious ways) from connectionist networks and therefore do not distinguish between the theories. Although we do not have space to review the entire literature here, consider two additional examples, frequency effects and double dissociations. We know from extensive prior research that high and low frequency verbs behave quite differently. Figure 1 presents the data from a past tense generation study reported by Seidenberg (1992) that examined high and low frequency regular and irregular verbs. There was an interaction between frequency and regularity: for high frequency items, there was no difference between regulars and irregulars but for lower frequency items, the irregulars took longer than the regulars (Prasada et al. 1990 presented similar results). The same interaction has been observed in studies of word pronunciation (Seidenberg et al. 1984; Taraban & McClelland 1988).

Insert Figure 1 About Here

One way to view these data is that frequency has different effects on rule-governed forms and exceptions--another dissociation taken as evidence for separate mechanisms. According to Pinker (1991) the 'associative net' is affected by frequency but rule-application is not. Again, however, such effects are known to be produced by connectionist networks for quite different reasons. Figure 1 presents the results from Daugherty and Seidenberg's (1992) model of past tense generation using the same types of stimuli as in the Seidenberg (1992) experiment. In brief, the explanation for this result is as follows. The model performs as well on high frequency irregulars as on regulars because repeated exposure to a pattern such as <u>take-took</u> washes out the effects of ‘enemies’ such as <u>bake-baked</u> and <u>fake-faked</u>. This occurs because the effects on the weights of exposure to the pattern itself are greater than the effects of exposure to neighboring items. Lower frequency exceptions yield poorer performance because they are presented less often and are more affected by exposure to rule-governed enemies. For further detail, see Plaut et al. 1996, who provide both a mathematical derivation of this interaction and simulation models that demonstrate it.

As a final example, consider the dramatic dissociations between regular and irregular forms that have recently been reported in brain-injured patients. Marslen-Wilson and Tyler (1997) presented data from studies of patients whose processing of verbs was impaired in specific ways. The patients exhibited a double dissociation: poorer performance on regulars compared to exceptions or poorer performance on exceptions compared to regulars. Such double dissociations are classically interpreted as evidence for separate neurological subsystems (Shallice 1988), in this case, the separate mechanisms in the Pinker theory. The validity of this form of reasoning from data to theory has been decisively challenged by connectionist modeling, which shows that instead of arising from selective damage to different subsystems, they can be produced by different types of damage to the single network mechanism.

The evidence supporting this assertion is provided by research on forms of dyslexia that are acquired as a consequence of brain injury (Plaut et al. 1996). Patients often show partial impairments in their ability to read words and nonwords aloud. Two prominent patterns of impairment have been termed 'surface' and 'phonological' dyslexia (Patterson 1981). Surface dyslexics are impaired at reading words with irregular pronunciations such as pint or brooch but their ability to pronounce 'rule-governed' items such as must and nonwords such as nust is relatively preserved. Phonological dyslexics show the opposite pattern: they are impaired at using the rules to pronounce nonwords and better at reading familiar words. Marlsen-Wilson and Tyler (1997) have now shown that similar impairments occur in use of the past tense. Like the impairments they reported, the dyslexic cases are standardly interpreted within a dual-mechanism theory (Coltheart et al. 1993). Surface dyslexics are said to have an impairment in the 'associative net' used in pronouncing irregular words but retain use of pronunciation rules; phonological dyslexics have an impairment in using the rules but retain use of the associative net.

Plaut et al. (1996) discussed how such effects arise from different types of damage to their network model. In their simulations, the phonological dyslexic data pattern derived from damage to the representation of phonological information, which has a bigger impact on nonword generalization than on familiar words. The surface dyslexic pattern was created by introducing a different form of damage (removing the contributions of semantic information), which has a bigger impact on exceptions than generalization. Why these outcomes obtained is discussed in Plaut et al. (1996) and Harm and Seidenberg (1997). Here the main point is that the double dissociation that Marslen-Wilson and Tyler report with respect to the past tense has also been observed in reading aloud; in both domains the dissociations have been taken as evidence for dual-mechanism theories; in both domains, there is a competing single-mechanism connectionist theory; and studies of implemented networks show how such double dissociations can arise without separate subsystems. These analyses have strongly challenged assumptions about the interpretation of double dissociations that date from 19th century neurology.

In summary, the kinds of dissociations between regulars and irregulars that J et al. and others have investigated do not actually distinguish between the competing theories. The same questions we have raised concerning the interpretation of J et al.'s data arise in connection with other types of data that have been offered as support for the dual-mechanism theory.

7. Conclusions. Our critique of the J et al. study has focused on problems with the conception of the study (whether the dissociations they sought could discriminate between the theories), the design of the experiment (the blocking problem), and the interpretation of the data (whether different brain regions were activated for regular and irregular verbs). On this basis we suggest that the study does not provide decisive evidence bearing on the debate.

In closing we offer some suggestions for future research. Although we have been critical of the J et al. study, it is clear that this kind of research has the potential to provide more decisive evidence bearing on the past tense debate and other questions about the brain bases of language. Neuroimaging is a relatively new methodology and changing rapidly as the technology evolves and researchers learn from initial successes and failures. Our analysis of the limitations of the J et al. study suggests several directions for future research.

First, it will be necessary to examine more refined hypotheses that have the potential to discriminate between theories; merely looking for overall differences between groups of regular and irregular verbs is not sufficient. One obvious possibility is to examine the interaction between frequency and regularity. The single and dual mechanism theories predict different patterns of brain activation for these items. Our account suggests that higher frequency regular and irregular forms should produce similar patterns of brain activity, whereas lower frequency regulars and irregulars should differ. The dual-mechanism theory, on the other hand, provides no basis for expecting any subgroups of regulars and irregulars to activate the same brain areas. We make the further prediction that similar outcomes should obtain in both the past tense and word pronunciation domains because they are governed by the same principles.

Second, it is essential to avoid the presentation of stimuli in homogeneous blocks, which introduced so many problems in the J et al. study. This constraint on experimental design means that the standard PET methodology is probably not suitable for studying the past tense. Better methods are becoming available, however. Recent developments in fMRI technology permit images to be derived from single trials (Buckner et al. 1996). This obviates the need to record over a block of similar items.

Third, it is important to equate the experimental conditions with respect to potentially confounding factors such as difficulty or how deeply the stimuli have to be processed in order to perform a task. Differences between regular and irregular items in reaction time, brain activation patterns, or other dependent measures cannot be attributed to type of verb unless such confounds are controlled. Pretesting of the stimuli could be used to create groups of regular and irregular items that are equated for difficulty; randomly intermixing different types of stimuli insures that each verb must be processed to the same degree and discourages formulating strategies in response to expectations about the composition of the stimuli.

Finally, as in other psycholinguistic research, it is necessary to obtain converging evidence from multiple studies that test competing hypotheses. We have suggested one such experiment: examine each type of verb (e.g., regular, irregular) at different levels of difficulty (e.g., easy, hard). Whereas J et al. asked whether different types of verbs activate different brain regions, it is also important to determine whether a single type of verb will activate multiple regions. Studies that address whether other kinds of stimuli and tasks also activate the regions identified in studies of the past tense would provide another source of converging evidence. With appropriately designed and conducted experiments, brain imaging studies could provide a clear resolution to this debate.

References

Barinaga, Marcia. 1997. What makes brain neurons run? Science 276.196-198.

Berko, Jean. 1958. The child's learning of English morphology. Word 14.150-177.

Bloch, B. 1947. English verb inflection. Language 23.399-418.

Buckner, Randy L., Peter A. Bandettini, Kathleen M. Ocraven, Robert L. Savoy, Steven E. Petersen, Marcus E. Raichle, and Bruce R. Rosen. 1996. Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. Proceedings of the National Academy of Sciences USA 93.14,878-14,883.

Bybee, Joan. 1985. Morphology. Philadelphia: Johns Benjamins.

Bybee, Joan L., & C.L. Moder 1983. Morphological classes as natural categories. Language 59.251-270.

Bybee, Joan, & Slobin, Dan I. 1982. Rules and schemas in the development and use of the English past tense. Language 58.265-289.

Coltheart, Max, Brent Curtis, Paul Atkins & Michael Haller.1993. Models of reading aloud: Dual-route and parallel-distributed approaches. Psychological Review 100.589-608.

Craik, Fergus I.M., & Robert S. Lockhart. 1972. Levels of processing: A framework for memory research. Journal of Verbal Learning and Verbal Behavior 11.671-684.

Daugherty, Kim., & Mark S. Seidenberg. 1992. The Past Tense Revisited. Proceedings of the 14th Annual Meeting of the Cognitive Science Society. Hillsdale, NJ.: Erlbaum Associates.

Demb, Jonathan B., John E. Desmond, Anthony D. Wagner, Chandan J. Vaidya, Gary H. Glover, and John D.E. Gabrieli. 1995. Semantic encoding and retrieval in the left inferior prefrontal cortex: A functional MRI study of task difficulty and process specificity. Journal of Neuroscience 15. 5870-5878.

Elman, Jeffrey L.., Elizabeth A. Bates, Mark J. Johnson, Annette Karmiloff-Smith, Domenico Parisi, and Kim Plunkett. 1996. Rethinking innateness: A connectionist perspective on development. MIT Press: Cambridge, MA.

Forster, Kenneth I., & Carol Davis. 1984. Repetition priming and frequency attenuation in lexical access. Journal of Experimental Psychology: Learning, Memory, & Cognition 10.680-698.

Gabrieli, John D., John E. Desmond, Jonathan B. Demb, Anthony D. Wagner, Maria V. Stone, Chandan J. Vaidya, and Gary H. Glover. 1996. Functional magnetic resonance imaging of semantic memory processes in the frontal lobes. Psychological Science 7.278-283

Gluck, Mark A. & David E. Rumelhart (eds.) 1990. Neuroscience and connectionist theory. Hillsdale, N.J.: Erlbaum Associates

Goldman-Rakic, Patricia S. 1988. Topography of cognition: parallel distributed ndetworks in primate association cortex. Annual Review of Neuroscience 11.137-156.

Gordon, Barry. 1983. Lexical access and lexical decision: Mechanisms of frequency sensitivity. Journal of Verbal Learning and Verbal Behavior 22.24-44.

Halle, Morris & K.P. Mohanan. 1985. Segmental phonology of modern English. Linguistic Inquiry 16.57-116.

Hare, Mary, & Jeffrey L. Elman. 1995. Learning and morphological change. Cognition 56.61-98.

Harm, Michael W., & Mark S. Seidenberg. 1997. Phonology, reading and dyslexia: Insights from connectionist models. Manuscript submitted for publication.

Henderson, Leslie. 1982. Orthography and word recognition in reading. London: Academic Press.

Hoeffner, James H. 1997. Are rules a thing of the past? A single mechanism account of English past tense acquisition and processing. Unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA.

Jaeger, Jeri J., Alan H. Lockwood, David L. Kemmerer, Robert D. Van Valin, Brian W. Murphy, and Hanif G. Khalak. 1996. A positron emission tomography study of regular and irregular verb morphology in English. Language 72.451-497.

Just, Marcel A., Patricia A. Carpenter, Timothy A. Keller, William F. Eddy, and Keith R. Thulborn 1996. Brain activation modulated by sentence comprehension. Science 274.114-116.

MacWhinney, Brian, & Jared Leinbach. 1991. Implementations are not conceptualizations: Revising the verb learning model. Cognition 39.121-157.

Marslen-Wilson, William. D., & Lorraine K.Tyler. 1997. Dissociating types of mental computation. Nature 387.592-594.

Patterson, Karalyn E., 1981. Neuropsychological approaches to the study of reading. British Journal of Psychology 72.151-174.

Petersen, Steven E., Peter T. Fox, Michael I. Posner, Mark Mintun, and Marcus E. Raichle. 1988. Positron emission tomographic studies of the cortical anatomy of single-word processing. Nature 331.585-589.

Pinker, Steven. 1991. Rules of Language. Science 253.530-535.

Pinker, Steven & Prince, Alan. 1988 On language and connectionism: Analysis of a Parallel Distributed Processing model of language acquisition. Cognition, 28, 73-193.

Plaut, D. C. 1997. Double dissociation without modularity: Evidence from connectionist neuropsychology. Journal of Clinical and Experimental Neuropsychology 17.291-321.

Plaut, David C., James L. McClelland, Mark S. Seidenberg, Karalyn E. Patterson, K. 1996. Understanding normal and impaired word reading: Computational principles in quasiregular domains. Psychology Review 103.56-115.

Poeppel, David. 1996. A critical review of PET studies of phonological processing. Brain and Language 55.317-351.

Prasada, Sandeep, Steven Pinker, William Snyder,1990. Some evidence that irregular forms are retrieved from memory but regular forms are rule generated. Paper presented at the 31st annual meeting of the Psychonomic society November: New Orleans.

Rayner, Keith, and Alexander Pollatsek. 1994. The psychology of reading. Hillsdale, N.J.: Erlbaum Associates.

Rice, Mabel (ed.) 1996. Toward a Genetics of Language. Hillsdale, N.J.: Erlbaum Associates

Roskies, Anne.L. 1994. Mapping memory with positron emission tomography. Proceedings of the National Academy of Sciences USA 91.1989-1991.

Rumelhart, David E. & James L. McClelland. 1986. On learning the past tense of English verbs. Parallel Distributed Processing: Explorations in the microstructure of cognition. Vol. 2, ed. by David E. Rumelhart, James L. McClelland and the PDP research group, 216-271. Cambridge, MA: MIT Press.

Rumelhart, David E., James L. McClelland, and the PDP Research Group eds. 1986. Parallel Distributed Processing: Explorations in the microstructure of cognition. Vol. 1. Cambridge, MA: MIT Press.

Scarborough, Daniel L., Cortese, Charles, & Scarborough, Hollis S. 1977. Frequency and repetition effects in lexical memory. Journal of Experimental Psychology: Human Perception & Performance 3.1-17.

Seidenberg, Mark S. 1992 Connectionism without tears. Connectionism: Theory and Practice, ed. by Stephen Davis, 84-122. Oxford University Press.

Seidenberg, Mark S. 1995. Visual word recognition: An overview. Speech, language, and communication, ed. by Joanne L. Miller and Peter D. Eimas, 138-180. San Diego: Academic Press.

Seidenberg, Mark S., & Maggie Bruck. 1990. Consistency effects in the generation of past tense morphology. Paper presented at the 31st meeting of the Psychonomic Society November: New Orleans.

Seidenberg, Mark S., & James L. McClelland. 1989. A distributed, developmental model of word recognition and naming. Psychological Review 96.523-568.

Sergent, Justine. 1994. Brain-imaging studies of cognitive functions. Trends in Neurosciences 17.221-227.

Shallice, Tim. 1988. From neuropsychology to mental structure. Cambridge University Press, Cambridge: UK.

Taraban, Roman, and James L. McClelland. 1987. Conspiracy effects in word pronunciation. Journal of Memory & Language 26.608-631.

Ullman, Michael T, Suzanne Corkin, Marie Coppola, Gregory Hicock, John H. Growdon, Walter J. Koroshetz, and Steven Pinker. 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 9.266-276.

Footnotes

1 As a reviewer pointed out, most linguistic accounts of these phenomena have used the same principles to capture regular and irregular forms, in contrast to Pinker's theory, which treats them as categorically different. Our approach is more in keeping with the former approach. What the connectionist framework adds to traditional accounts is a way to represent regularities that vary in degree and type and a mechanism for acquiring much of this knowledge from examples.

2 We use the terms regular/rule-governed, irregular/exception and nonce/nonword interchangeably.

3 Readers who want to collect their own data will find that total time to generate past tenses for all of the words in a condition, measured to the nearest second, should yield reliable results. Note that although the regular and irregular stimuli have been matched in terms of initial phoneme, a factor that greatly affects spoken latencies, we have not made any attempt to equate them in terms of other factors such as frequency or length. This is a demonstration, not a formal experiment.

4 In reviewing this article, Jaeger challenged this account on the following grounds. All parties agree that generating the past tense of a verb like take requires recognizing the letter string as a specific lexical item, the one that is associated with the past tense took. Jaeger suggests that generating the past tense of a regular verb like bake also requires the word to be recognized, because otherwise the subject would not know which allomorph of the regular suffix to produce. Thus the fact that subjects knew in advance that every verb in the regular condition had a regular past tense would not afford the shallower response strategy we have described. However, the correct realization of the regular suffix depends only on the final segment of the present tense verb. All the strategy requires is that the subject be able to compute the phonological code of the verb; the literature on visual word recognition shows that good readers routinely and automatically compute this information (see Seidenberg, 1995, for review). This computation does not require lexical access; if it did, subjects would not be able to generate correct past tenses for nonce verbs like WUG, which have no representations in lexical memory to access. Seidenberg and McClelland (1989) and Plaut et al. (1996) describe implemented computational models of orthographic to

phonological conversion that apply without regard to the lexical status (word vs. nonword) of the letter string.

Both Jaeger and an anonymous reviewer have suggested that this strategy could only be achieved by using a rule, thereby providing evidence for the rule-based theory and against the connectionist account. Our assertion is that similarities across trials facilitated performance in the regular and nonce conditions. In connectionist networks, such effects can be implemented by assuming that each trial causes a small amount of updating of the weights. A large number of repetitions of particular patterns creates facilitation. The fact that the allomorphs overlap phonetically would also contribute to this effect. Thus, whether subjects used this strategy or not is orthogonal to whether they were using a rule or a connectionist net.

[5] Data similar to Seidenberg and Bruck's were reported by Prasada, Pinker, and Snyder (1990).

[6] This data set can be obtained from the present authors.

Table 1

Brain Areas Activated by Each Type of Stimulus

	Regular	Irregular	Nonce
1) Regular Only			
2) Irregular Only			
L. mid.temporal g.(21)		*****	
L. calcarine cortex (17)		*	
L. fusiform g. (19)		*	
3) Nonce only			
L. motor strip (4)			*
L. cuneus (18)			*
L. sup. occipital g. (19)			*
4) Regular/Irregular			
L. inf. parietal lob. (40)	*	**	
L. precuneus (mesial 7)	*	***	
R. precuneus (mesial 7)	*	*	
5) Regular/Nonce			
L. dorsolat. prefrontal (46)	*****		*****
L. ant. cingulate g. (24)	*		*
6) Irregular/Nonce			
L sup. frontal g. (10)		*******	*****
L. sup. parietal lob. (7)		**	*
L. inf. occipital g. (18)		*	*
Cerebellum			
R. hemisphere		**	*
L. hemisphere		**	*

midline		**	*
Other regions			
colliculi		*	*
7) All three: Regular, Irregular, Nonce			
L mid. frontal g. (44,45)	*	**	*
L. inf. frontal g. (44, 45)	*	***	**

Note: This table is derived from Jaeger et al.'s (1996) Table 4. The asterisks indicate the significance level of the PET activation for each area in each generate past-read verb comparison (i.e., Regular Past-Read Verb, Irregular Past-Read Verb, and Nonce Past-Read Nonce). One asterisk=omnibus, p<0.01, two asterisks=omnibus, p<0.001, and three asterisks indicate significance at Bonferroni, p<0.05 level.

Table 2

Irregular and Regular Verbs Blocked by Type

IRREGULAR		REGULAR	
fall	eat	flip	ease
build	bring	bill	brush
shoot	take	shock	brush
dig	lose	dim	last
sit	stand	sip	stake
spend	sleep	spell	sleet
sweep	catch	swat	cart
feed	come	fear	cope
wear	speak	want	sport
break	hide	braid	hope
hear	fly	heat	fry
blow		bloat	

Table 3

Irregular Verbs Grouped by Subtype

1. sweep	13. cling
2. sleep	14. fling
3. weep	15. wring
4. keep	16. string
5. leap	17. swing
6. bend	18. cut
7. spend	19. shut
8. send	20. let
	21. set
9. blow	22. hurt
10. know	23. hit
11. grow	
12. throw	

Table 4

Results of Jaeger et al. (1996) and Seidenberg (1992) Studies

	JAEGER ET AL.		SEIDENBERG	
	Correct	*Net*	*Correct*	*Net*
Read Verb	486	-	575	-
Regular Past	524	38	925	350
Irregular Past	724	238	1030	455
Read Nonce	628	-	650	-
Nonce Past	684	56	1312	662

Note: Data are response latencies in milliseconds for correct trials. Jaeger et al. also provided data for all trials that yielded similar results. The Net data are for the Past - Read subtractions.

Figure Caption

Figure 1. Left: data from a behavioral study in which subjects generated past tenses for verbs. Net generation effects = latency to generate past tense - latency to name present tense verb aloud. See Seidenberg (1992) for details. Right: data from a simulation model using the same kinds of stimuli. See Daugherty and Seidenberg (1992) for details. Labels on the y-axis have been multiplied by 100 for legibility.

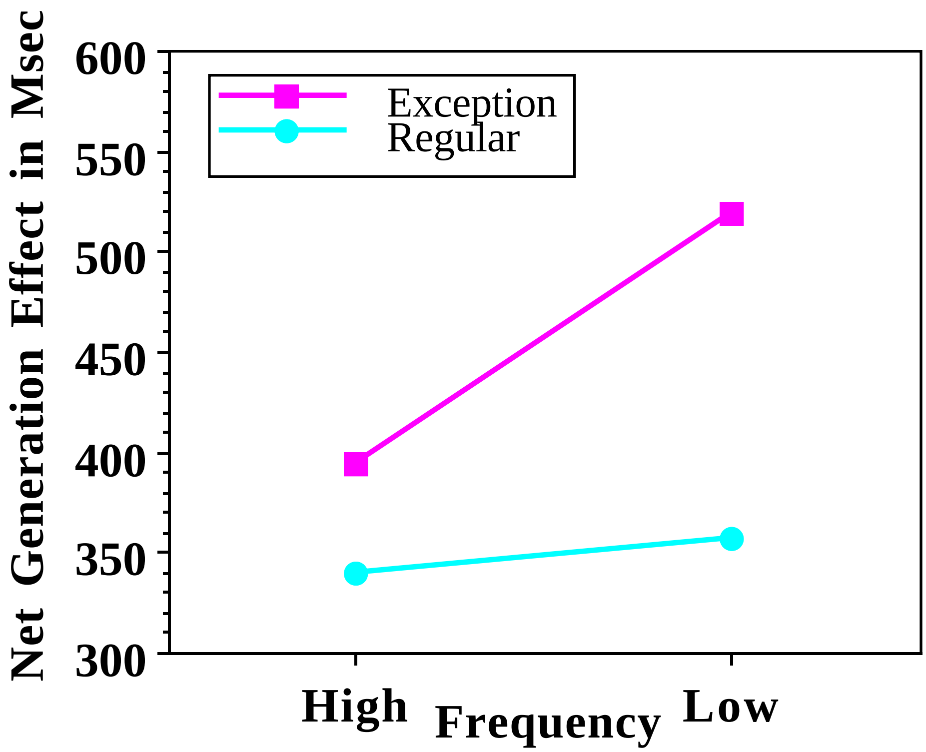
Net Generation Effect in Msec
600
550
500
450
400
350
300
Exception
Regular
High
Frequency
Low

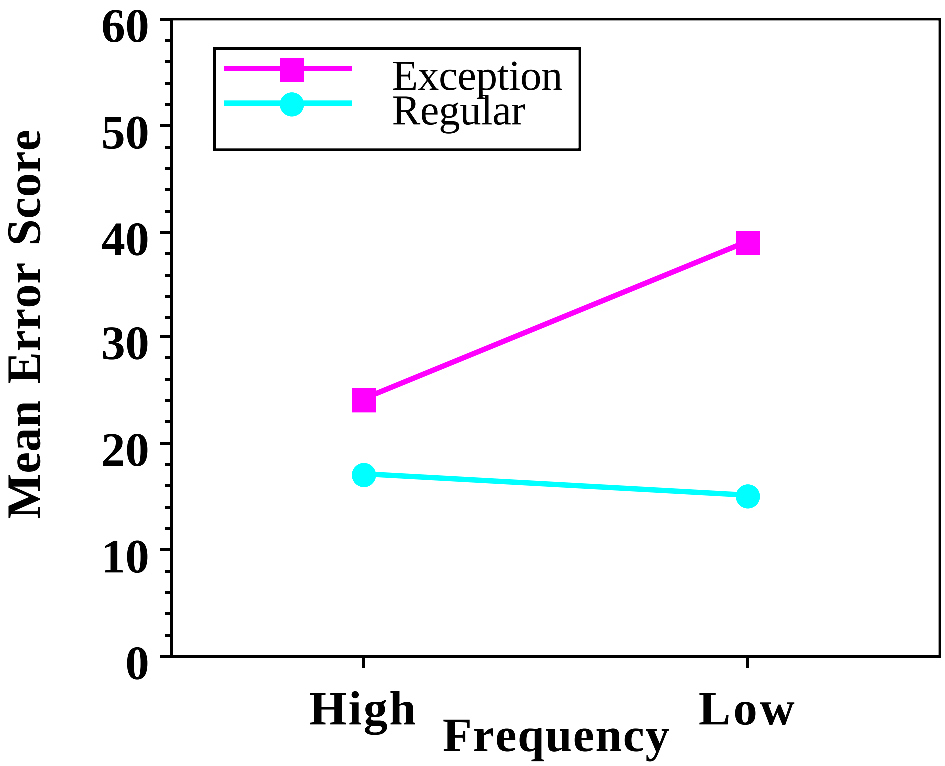
Mean Error Score
60
50
40
30
20
10
0
Exception
Regular
High
Frequency
Low