Journal of Psycholinguistic Research, Vol. 27, No. 1, 1998

Lexical Representation of Gender: A Quasiregular Domain

Marcus Taft[1,3] and Fanny Meunier[2]

Accepted May 27, 1997

The relationship between gender and word ending in French is a quasiregular one (e.g., most words ending in -ette *are feminine, but not all). As such, the gender of low-frequency irregular forms (e.g.,* squelette, *which is masculine) should take longer to classify than low-frequency regular forms according to neural network models. A regularity effect was found in Experiment 1, but it did not interact with word frequency. It was further revealed that there was difficulty in making gender decisions (Experiment 2) and gender verification responses (Experiment 3) to words whose endings were highly informative of gender, but whose associated article was not. These words were place names beginning with a vowel, like* Australie, *which do not take an indefinite article and whose definite article is ambiguous (l'). How a neural network might handle these results is discussed, and an alternative account is considered whereby there are two potential sources of gender information, lexical and nonlexical, with the latter being used to confirm the former.*

A problem that is familiar to most English-speaking students learning French is the difficulty encountered in remembering the gender of nouns. It is hard to recall whether the noun is masculine and therefore takes the definite article *le* and the indefinite article *un*, or whether it is feminine and therefore takes the definite article *la* and indefinite article *une*. From their meanings, a word like *femme* (*woman*) is obviously feminine and *oncle* (*uncle*) is obviously masculine, but what about *table* and *menu*? The meaning of these words suggests nothing about their gender, yet one still needs to know their

[1] University of New South Wales, Sydney NSW 2052, Australia.
[2] Université René Descartes, 75006 Paris, France.
[3] Address all correspondence to Marcus Taft, School of Psychology, University of New South Wales, Sydney NSW 2052, Australia.

0090-6905/98/0100-0023$15.00/0

gender in order to match them with the appropriate form of the article or adjective [e.g., *belle* (*beautiful*) for feminine nouns or *beau* (*beautiful*) for masculine nouns]. The focus of the present research is how this "grammatical gender" is represented in the French speaker's lexical memory.

Of course, many other languages have grammatical gender as well, but in some it at least seems that the ending of the word helps to discriminate between the different genders. For example, in Spanish, final *-o* is almost always masculine and final *-a* is almost always feminine. In French, however, the relationship between form and gender appears to be arbitrary, at least to the novice. In fact, it is not arbitrary. There are many word-endings in French which are highly predictive of gender. For example, 98% of nouns ending in *-ie* are feminine (e.g., *vie*, *folie*), and 99% of nouns ending in *-t* are masculine (e.g., *tricot*, *incident*). There exist a number of such highly predictive endings (e.g., *-ette*, *-ue*, and *-tion* are over 95% predictive of being feminine, and *-i*, *-u*, and *-age* of being masculine), though there are also endings which are not so predictive (e.g., *-le* following a consonant, as in *triangle*, *cercle*, and *table*, indicates masculine gender only about 70% of the time). Note that the orthographic form of the ending is more indicative of gender than is its phonological form given that homophonic endings (e.g., *-ue* and *-u*) can be associated with different genders.

Tucker, Lambert, Rigault and Segalowitz (1968) were probably the first psychologists to pay attention to the relationship between gender and word-ending. They presented native French speakers with nonexistent words that had a range of typical endings and asked them to assign gender to these items. What they found was that the gender given to a nonword was correlated with the proportion of times its ending was actually associated with the masculine or feminine gender in French. From this and other similar research, Tucker, Lambert, and Rigault (1977) concluded that native French speakers have synthesized a set of rules which associate word-endings with grammatical gender. Assignment of gender was seen as "an example of rule-governed behavior" (Tucker et al., 1977, p. 64).

One cannot really say, though, that the relationship between orthography and gender is regular: Some endings are highly predictive of one gender, but others are not. There is a range of predictiveness. What this means is that the relationship is a quasiregular one, that is, there is considerable regularity, but exceptions exist. Now, it is in such quasiregular domains that neural network models like those proposed by Seidenberg and McClelland (1989) and Plaut, McClelland, Seidenberg, and Patterson (1996) have excelled in simulating patterns of human performance. The system is trained on a set of input-output correspondences, including both regular and irregular ones, and from this a set of input units becomes linked (via a layer of hidden units) to a set of output units in a particular way, with a range of

different link-strengths depending on the frequency of the particular relationship. When a new input is given, the system uses these links to activate an appropriate output.

The quasiregular domain that has been the primary focus of these models has been the relationship between orthography and phonology: *Grand* is a regular word because each of its graphemes (*g*, *r*, *a*, *n*, and *d*) are given their most typical pronunciation in this word (/g/, /r/, /æ/, /n/, and /d/, respectively) whereas *great* is irregular because the grapheme *ea* is not given its most typical pronunciation in this word (i.e., /eɪ/ instead of /i:/). When English readers are asked to name such words aloud, it is typically found that irregular words take longer to name than regular words, but only for lower-frequency words (e.g., Andrews, 1982; Seidenberg, Waters, Barnes, & Tanenhaus, 1984), and this is exactly the pattern of output arising from simulations using the neural network models (see Seidenberg & McClelland, 1989).

Another aspect of the relationship between orthography and phonology that has been found to have an impact on naming responses is the "consistency" of the pronunciation of subword units (e.g., Andrews, 1982; Glushko, 1979; Seidenberg et al., 1984; Stanhope & Parkin, 1987). A regular word like *rush* takes longer to name than a regular word like *mill* because the subword unit *-ush* has more than one pronunciation (*rush*, *hush*, *mush* on the one hand and *bush*, *push* on the other), while *-ill* has only one. Again, such an effect interacts with word frequency (Seidenberg et al., 1984) and has been simulated by neural network models (Plaut et al., 1996; Seidenberg & McClelland, 1989).

The interaction of both the regularity and consistency effects with word frequency arises within the neural network from the fact that the activation relationship between input and output is curvilinear, with a function that asymptotes at the higher levels (see Plaut et al., 1996, p. 74, Fig. 7). Thus, any variations in the input characteristics of highly activated words (i.e., words which have been frequently exposed during training) have less impact on output than those of more weakly activated words. In this way, variations in the regularity and consistency of low-frequency words will have more impact on the responses to those words than will such variations for high-frequency words.

Now, one of the features of a neural network model is the fact that it simply makes use of statistical relationships that exist between different aspects of the domain in which it is working, and is indifferent to what the actual content of that domain is. That is, any input and any output could be looked at as long as there is a relationship that exists between them. As Plaut et al. (1996) wrote, the neural network approach "is applicable, we believe, to a wide range of linguistic and cognitive domains—essentially to

all those with a quasi-regular structure . . .'' (p. 106). Therefore, while neural networks have been primarily directed toward cases where the input is a graphemic representation and the output a phonemic one, they could alternatively be directed at cases where the input is a graphemic (or phonemic) representation of a noun and the output is a description of its gender. In fact, MacWhinney, Leinbach, Taraban, and McDonald (1989) have looked at exactly this situation when they demonstrated that a network could be trained to accurately assign the appropriate article to a German noun primarily on the basis of phonological cues (their Model 3). The reliability of a phonological cue was shown to be a factor that affected the acquisition of gender, that is, when the phonological information was consistently associated with a particular gender, the association was readily learnt by the network. This is equivalent to a ''gender consistency effect.'' What the work of MacWhinney et al. does not tell us, however, is whether adult language users show such a consistency effect in generating the gender of a noun and, furthermore, whether the characteristics of the relationship between gender and phonology (or orthography) follow the same sort of pattern as has been observed when adult language users respond in other quasiregular domains, in particular, the domain of grapheme-phoneme relationships. The present research aimed to explore this latter issue.

The consistency and regularity effects that are observed in the domain of grapheme-phoneme relationships are measured by looking at naming responses to visually presented stimuli. This task takes graphemic input and taps directly into the phonemic output. In order to look at gender consistency and regularity effects, then, we need a task that takes graphemic (or phonemic) input and taps directly into the gender output. Such a task is the ''gender decision task'' whereby subjects are presented with a noun and must respond whether it is masculine or feminine (a task first used by Radeau, Mousty, & Bertelson, 1989). If the quasiregular relationship between orthography and gender behaves in the same way as that between orthography and phonology, there should be gender regularity and consistency effects when making gender decisions, but only for lower-frequency words as with the phonological regularity and consistency effects.

In fact, the gender consistency effect has already been examined. Desrochers, Paivio, and Desrochers (1989) found faster gender decision responses to visually presented words whose ending was predictive of gender (e.g., *tricot* being masculine) than to words whose ending was not so predictive (e.g., *triangle* being masculine). ''Predictiveness'' is essentially equivalent to ''consistency'' in that predictive endings are ones which are consistently associated with one gender while nonpredictive endings are ones which can be associated with either. The nonpredictive items used by Desrochers et al. were always nouns which had the most typical gender for

that ending (e.g., there are more masculine nouns ending in *-le* following a consonant than feminine ones) and, in this sense, they can be seen as being regular words despite being inconsistent. So, the finding that *triangle* takes longer than *tricot* to identify as being masculine is equivalent to the finding that *rush* takes longer to name than *mill.*

Interestingly, although there was an interaction between the predictiveness effect and word frequency on error rates, the size of the predictiveness effect on response times was just as great for high-frequency words as it was for low-frequency words. Desrochers et al. (1989) put more weight on this lack of interaction than on the error data and, therefore, gave a postlexical interpretation, that is, the predictiveness of the word-ending plays its role after the representation of the word has been accessed in lexical memory. The nature of the word-ending is one of the pieces of information that is activated in the process of deciding what the gender of the word is, but the full details of how the decision is reached are unclear.

A postlexical account of gender assignment was also given by Bates, Devescovi, Pizzamiglio, D'Amico, and Hernandez (1995) on the basis of a study using auditory presentation of Italian words. Their correlational analysis showed an effect of the predictiveness of the word ending (which they called "transparency") in a gender decision task (which they called "gender monitoring"), but not in a word repetition task. They suggested that the effect of word-ending in gender judgments comes into play during a postlexical checking procedure which is not required when merely repeating the word. Surprisingly, Bates et al. failed to find any effect of word frequency in the gender decision task, but they put this down to the fact that the task was an auditory one where frequency effects are hard to obtain.

The first experiment to be reported here looked for a gender regularity effect in visually presented French nouns of high and low frequency and, as such, complemented the research of Desrochers et al. (1989) looking at consistency.

EXPERIMENT 1

Irregular gender refers to the situation where the gender of a noun contradicts the gender that would be expected from the word-ending. For example, *squelette* is a masculine noun, which makes it irregular because only 1% of words ending in *-ette* are masculine. On the other hand, *socialisme* is also masculine, but is regular because all words ending in *-isme* are masculine. Similarly, *vertu* and *figure* are both feminine, but the former is irregular and the latter regular because words ending in *-u* are feminine in only 2% of cases while words ending in *-ure* are feminine 94% of the time.

The experiment required subjects to decide the gender of a mixed set of regular and irregular nouns. If orthographic information contributes to gender decision, as the results of Desrochers et al. (1989) suggest, then the irregular words should take longer to judge than the regular ones. Furthermore, according to a neural network account, this regularity effect should be greater for low-frequency words than for high-frequency words.

Method

Subjects

Sixteen native French speakers from Université René Descartes, Paris, served as subjects in the experiment.

Materials and Procedure

Subjects were presented with 40 masculine nouns and 40 feminine nouns randomly mixed together for a gender decision to be made via the pressing of a button. Each word was presented without any article. The items were designed in pairs such that one member of a pair had an ending that was highly typical of its gender (e.g., *folie*, which is feminine and has the feminine ending *-ie*; *incident*, which is masculine and has the masculine ending *-t*). These were the "regular" words. The other member of the pair was "irregular," having a highly unusual ending for its gender (e.g., *forêt* is feminine though it ends in *-t*; *incendie* is masculine though it ends in *-ie*). The members of each pair were matched on gender, on word frequency (according to *Trésor de la Langue Française*, 1971) and on approximate length, and half of the pairs were high-frequency words (>25 per million; average of 104 per million[4]) while half were low-frequency words (<12 per million; average of 4.3 per million). Examples of high-frequency words are *vertu*, *incendie*, *forêt*, *monde* (irregular), and *figure*, *incident*, *folie*, *jour* (regular) while examples of low-frequency words are *glu*, *squelette*, *perdrix*, *mercure* (irregular) and *taie*, *socialisme*, *palette*, *museau* (regular).

Items were displayed for as long as it took the subject to respond, at which point an asterisk appeared on the screen for 100 ms, acting as the fixation point for the next item. There were 10 practice items. Subjects were instructed to provide the gender of each presented word by means of button press. They were told to respond as quickly but as accurately as possible, and that if they could not think of the gender they should make a guess.

[4] The frequency presented is the antilog of the mean log frequency.

Results

The gender decision times and error rates are presented in Table I. Although there was no reason to expect any difference between the two types of response ("masculine" and "feminine"), this gender factor was nevertheless included in the analysis. A significant frequency effect was revealed on reaction time, $F_1(1, 15) = 20.50, p < .001$; $F_2(1, 37) = 4.16$, $p < .05$, but not on error rates, $F_1(1, 15) = 1.67, p > .1$; $F_2(1, 37) = 0.49$, $p > .1$, along with a significant effect of regularity on both reaction time, $F_1(1, 15) = 27.22, p < .001$; $F_2(1, 37) = 7.50, p < .01$, and errors, $F_1(1, 15) = 8.71, p < .01$; $F_2(1, 37) = 19.68, p < .001$. There was no sign of an interaction between frequency and regularity on either reaction time or error rates, all Fs < 1.

The main effect of gender was not significant, $F_1(1, 15) = 3.99, p > .1$; $F_2(1, 37) = 2.89, p > 1$. The size of the regularity effect for reaction times was greater for feminine than masculine nouns, but only on the subject analysis, $F_1(1, 15) = 5.32, p < .05$; $F_2(1, 37) = 0.22, p > .05$, while the reverse was true for error rates on the subject analysis, $F_1(1, 15) = 6.82, p < .05$; $F_2(1, 37) = 3.15, p > .05$, suggesting that there is no genuine effect of gender category in the experiment. Neither the interaction of gender and frequency, nor the three-way interaction of gender, frequency, and regularity, was significant, though the latter came close on the subject analysis of errors, $F_1(1, 15) = 3.46, p > .05$.

Discussion

These results clearly showed that the gender of a noun was generated in a way that was influenced by orthographic information at the end of the word. When that ending was rarely associated with the gender of the word, it was relatively difficult to identify that gender. Neural network accounts predict this regularity effect because their functioning centers on the statistical relationships between input (orthographic information) and output (gender identity). However, contrary to what was predicted by analogy to the quasiregular domain of pronunciational regularity, the results suggest that orthographic information was used in a way that is unaffected by word frequency.

One interpretation of the lack of interaction between regularity and frequency is that the word was accessed (hence showing frequency effects) and then orthographic information was used as an aid in determining gender at this point. This is the interpretation given by Desrochers et al. (1989) for their equivalent results in relation to the predictiveness of orthographic endings. The problem with this, however, is that it is hard to see why the impact

Table I. Gender Decision Times (Reaction Time, or RT, in Milliseconds) and Percentage Error Rates (ER) for Experiment 1; Standard Error is in Parentheses

	Regular			Irregular		
	Example	RT	ER	Example	RT	ER
		Masculine				
Low frequency	socialisme	797	2.5%	squelette	852	15.0%
		(51)	(.14)		(45)	(.30)
High frequency	incident	747	1.9%	incendie	784	10.6%
		(33)	(.10)		(37)	(.34)
		Feminine				
Low frequency	palette	742	4.4%	perdrix	819	7.5%
		(39)	(.21)		(37)	(.30)
High frequency	figure	686	2.5%	vertu	776	8.8%
		(36)	(.11)		(44)	(.35)

of the orthographic ending must rely on accessed lexical information. If one is basing one's gender decision on orthographic cues, as the regularity and consistency effects suggest, then one should be able to achieve this without actually accessing the word in lexical memory. Tucker et al. (1968, 1977) demonstrated that this is possible when they showed that the gender that subjects assigned to nonwords was a reflection of the predictiveness of orthographic ending. That is, orthographic cues can be used even when no lexical information is accessible.

When the gender judgment is made to real words, however, it is clear that lexical information *is* being accessed. Otherwise, there would be no word frequency effect. So, it seems that while lexical information is not *required* for the use of orthographic cues, both lexical information and orthographic cues contribute to the judgment of gender. Perhaps there are two sources of gender information: a "lexical" source which consults information stored with the word (e.g., "*folie* is feminine") and another "nonlexical" source which consults orthographic information (e.g., "*-ie* is a feminine ending"). The response must ultimately be based on the information derived from the lexical source if the subject is going to accurately categorize the gender of irregular and nonpredictive words, but the nonlexical source can influence this decision (perhaps during a final checking stage, as suggested by Bates et al., 1995). Regularity effects arise when the output of the two sources conflict, while consistency effects arise when orthography is not helpful (e.g., *-le* is an unreliable guide to gender). Frequency and regularity do not interact because the nonlexical orthographic information is only used once the lexical information becomes available.

What is the nature of the lexical information that is used in determining gender? The implicit assumption we have made so far is that it takes the form of some sort of abstract label associated with the noun, either "masculine" or "feminine". It is debatable, however, if nouns are ever specifically given such labels since these are never required in the normal course of language production and reception. Knowledge of grammatical gender is only relevant to the French speaker in so far as it is usually required for the correct selection of the article to go with the noun (*une* or *un*; *la* or *le*), the correct selection of any anaphoric pronouns (e.g., *elle* or *il*), and also of the appropriate form of the adjective if there is one (e.g., *belle* or *beau*; *froid* or *froide*). What this suggests is that lexical knowledge of gender at least includes information about the article. Such a claim is supported by Deutsch and Wijnen (1985) from their study of Dutch gender, and also by the finding in French that gender decision responses are faster when subjects respond with the labels *un* or *une* compared to *maculin* or *féminin* (Desrochers et al., 1989). In addition, we often observed subjects making their response after overtly uttering the article (usually the indefinite article) and word together before responding.

According to the "dual-source" explanation, nonlexical orthographic information is only used once the lexical information about gender becomes available, and we can test this by seeing whether gender judgments become difficult when lexical information is impoverished even though orthographic cues are available. In particular, if we assume that the article makes a major contribution to lexical gender information, we can look specifically at words for which the article cannot be used to help determine gender. Under what circumstances does this happen?

Names of places which are larger than a city take the definite article in French, but not the indefinite article. So, one can have *le Michigan*, *la Russie*, and *le Brésil*, but not *un Michigan*, *une Russie*, or *un Brésil*.[5] In addition, when the place-name begins with a vowel, the definite article itself reduces to *l'* and therefore cannot be used to distinguish masculine and feminine as in, for example, *l'Oregon*, *l'Australie*, and *l'Israël*. Thus, information about articles cannot be used to decide that *Israël* is masculine. However, orthographic information can. In particular, with very few exceptions (e.g., *le Zaïre*, *le Tennesse*), all place-names that end in *-e* are feminine and all the rest are masculine. Even ordinary words are overwhelmingly masculine (99%) when ending in *-l* and therefore *l'Israël* is highly likely to be masculine, and the same is true for the *-on* ending of *Oregon*,[6] while the

[5] In fact, indefinite articles are occasionally used with place-names in a journalistic sense, as in *il pleut sur une France devastée*, but this use is very rare.

[6] The ending *-on* is meant to refer only to those cases where it follows a consonant, since *-ion* is a predominantly feminine ending.

-ie ending of *Australie* is almost inevitably feminine. So a second experiment was set up whereby place-names with highly predictive endings (which is most of them) were presented for gender decision in order to determine the relative importance of article information and orthographic information in generating the gender of a word.

Experiment 2

In Experiment 2, place-names were presented without any article and half began with a consonant (e.g., *Russie*, *Brésil*) and half with a vowel (e.g., *Australie*, *Israël*). All had highly predictive endings. According to the dual-source idea, there should be considerable difficulty responding to the names beginning with vowels. The nonlexical source cannot be used to determine the gender on the basis of the orthographic ending until gender information is extracted from lexical knowledge, and gender is therefore difficult to determine when the lexically stored article is not useful in providing gender information (i.e., when it only takes the form *l'*).

What the neural network account predicts depends on the nature of the output units. If it is assumed that they take the form of the abstract labels "masculine" and "feminine," then it need not be the case that place-names beginning with vowels will cause any difficulty when their orthographic endings are highly predictive of their gender. The regularity effect arises from the exposure of the system to regularity of association between orthographic ending and gender which is captured during training and used by the system in determining gender. Other information associated with a word during training will develop links to gender output as well, as long as it is systematically associated with the gender of the word. The article that typically occurs in conjunction with a word is an example of such information.[7] Now if the article is uninformative about gender, as in the case of place-names beginning with a vowel, then the primary source of activation to gender output will be the orthographic cues. That is, the network extracts systematic associations between input and output no matter what their nature is, and therefore, any systematic associations will contribute to the output. In this way, place-names beginning with a vowel should cause no particular difficulties in making gender judgments because their predictive orthographic endings will provide the gender information required.

[7] Word meaning constitutes another type of information that may be systematically associated with gender (see MacWhinney et al., 1989). However, this factor was controlled in the present experiment by using place-names which have no systematic semantic association with gender.

Alternatively, however, the output units might not take the form of gender labels, but may directly represent articles. This is how the network designed by MacWhinney et al. (1989) is set up. What the network learns is an association between orthographic (or phonological) input and the articles that are found in conjunction with that orthographic form. The gender of a noun is then determined from the gender of the outputted article. In this way, when the input is *Russie*, the output is *la* and a "feminine" judgment can therefore be made. But what about *Australie*? The only article associated with *Australie* is *l'* and, because this is ambiguous in relation to gender, there should be difficulty in making the gender judgment.

So, if gender decisions about place-names beginning with vowels are more difficult than those about place-names beginning with consonants, it would suggest either that the dual-source account is correct or that the output units of a neural network represent the article. On the other hand, if there is no difference between the two types of place-name, it would support a neural network approach whose output takes the form of gender labels.

Method

Subjects

The same 16 subjects who were tested in Experiment 1 were also tested in Experiment 2.

Materials and Procedure

Eighteen pairs of place-names were constructed where one member of the pair began with a vowel and the other a consonant, but where the two names had the same ending (e.g., *Israël* and *Brésil*, *Uruguay* and *Paraguay*, *Islande* and *Thaïlande*). Half of the items were masculine and half were feminine, and they were all presented without any article.

The names were also matched on familiarity ratings collected from 20 independent French speakers. These raters were given a list of 112 place-names and instructed to rate them for familiarity on a 5-point scale, with a rating of 1 meaning *unknown* and 5 meaning *very well known*. Raters were told to base their judgments of familiarity on the name of the place rather than on the place itself. Two orders of list were used. The mean rating of the 18 item pairs selected from these lists was 3.65 for each of the two conditions.

The experimental procedure was the same as in Experiment 1.

Results

Gender decision times and error rates are found in Table II. Analysis revealed a highly significant difference between the place-names beginning

Table II. Gender Decision Times (Reaction Time, or RT, in Milliseconds) and Percentage Error Rates (ER) for Experiment 2; Standard Error is in Parentheses

	Vowel onset			Consonant onset		
	Example	RT	ER	Example	RT	ER
Masculine	Israël	1112	21.3%	Brésil	895	6.3%
		(104)	(1.21)		(56)	(.23)
Feminine	Australie	871	10.0%	Russie	719	2.5%
		(55)	(.52)		(19)	(.13)

with vowels and those beginning with consonants on both reaction time, $F_1(1, 15) = 10.59, p < .01$; $F_2(1, 17) = 35.28, p < .001$, and error rate, $F_1(1, 15) = 8.05, p < .05$; $F_2(1, 17) = 22.36, p < .001$.

In addition, judgments of masculine place-names took significantly longer than those of feminine place-names, $F_1(1, 15) = 15.41, p < .001$; $F_2(1, 16) = 16.99, p < .001$, with no interaction between gender and onset type, $F_1(1, 15) = 0.84, p > .05$; $F_2(1, 16) = 2.15, p > .05$. The same was essentially true for errors, where masculine cases were harder than feminine cases at least on the item analysis, $F_1(1, 15) = 3.65, p > .05$; $F_2(1, 16) = 7.28, p < .05$, with no interaction between gender and onset type, $F_1(1, 15) = 0.81, p > .05$; $F_2(1, 16) = 3.09, p > .05$.

Discussion

It seems that it was very difficult for French speakers to determine the gender of words for which the article provided no information. They ultimately performed at a level much better than chance, so this means that the gender information was eventually extracted, but they were particularly slow in so doing.

From a *post hoc* analysis, it was further observed that the familiarity of the place-name significantly correlated with response time ($r = -.70$ for vowel and $-.93$ for consonant onsets), but did not correlate with the size of the vowel/consonant response time difference ($r = +.17$). That is, there was no interaction between familiarity of a place-name (i.e., the subjective version of its frequency) and informativeness of the article that went with that place-name, and this suggests that gender information was extracted after the lexical representation of the word was accessed. Thus, it seems that familiarity influenced the speed of access to the lexical representation of a word, and the gender decision was made at a constant time after this, but that this constant value was considerably greater when the article associated

with the word was not informative about gender (i.e., when the name began with a vowel).

Such a conclusion is consistent with both the dual-source idea and a neural network whose output units represent the article (and any other information that is systematically found in association with the word). In the former case, lexical information about gender (e.g., the associated article) is extracted and then rules are used, perhaps to confirm the gender classification. Therefore, when the lexical information about gender is impoverished because the only available article is ambiguous (i.e., *l'*), it is difficult to make use of the rule based information (i.e., the fact that the word has a gender-predictive ending) and a gender decision is therefore hard to make. According to the neural network account, the predictiveness of the ending plays a role in generating the article as output, but if the only article so generated is ambiguous then the gender decision cannot be based upon that output.

It is then necessary for both accounts to explain how the correct classification of gender is eventually made when it cannot be based on the informativeness of the article. The dual-source account would need to say either that the predictive orthographic endings are eventually utilized or that other lexical information about gender is used, like knowledge of the gender of adjectives that can go with that place-name (e.g., feminine *belle* or masculine *beau*) or of its pronominal form (e.g., feminine *elle* or masculine *il*). The neural network account gives a neater explanation for how this other gender information might come into play. That is, during the course of exposure to the place-name in meaningful contexts, it might well become associated with adjectives and pronouns. When these betray the gender of the place-name (e.g., *elle est belle*, *l'Australie*), an association builds up not only between *Australie* and its article *l'*, but also between *Australie* and the adjective *belle* and the pronoun *elle*. If these also form output units along with the article, then they can be used to determine the gender.

The finding that "feminine" judgments are easier to make than "masculine" judgments is a potentially interesting one because it might indicate that responses are affected by the fact that the "feminine rule" is more defined than the "masculine rule." That is, feminine gender is associated with final *-e* only, while masculine gender is associated with all other endings. However, pursuing this line of thought is compromised in this experiment by the fact that it may simply be a familiarity effect. It turns out that the familiarity ratings for the feminine place-names (mean of 4.24) were significantly higher than those for the masculine place-names (mean of 3.08), $t(19) = 3.57, p < .01$.

Now, the gender decision task requires the reader to bring the gender of the word into consciousness in order to respond. As mentioned earlier, it

is rarely the case that the reader or speaker has to state the gender of a noun explicitly. Instead, the way that gender normally manifests itself is in establishing agreement between a noun and its article, adjective, or anaphoric pronoun. For this reason, the gender decision task could be criticized as being open to influences and strategies which are specific to that task and as therefore not reflecting the normal processing of gender.

In the third experiment, then, a different task was used whereby the subject was not explicitly required to classify the gender of the noun. Instead, an adjective was paired with the place-name and the task was to say whether the form of the adjective was appropriate. This was a "gender verification" task. For example, a *yes* response should be given to *Australie profonde* since the the final *-e* of the adjective *profonde* (*deep*) is appropriate for a feminine word like *Australie*, whereas a *no* response should be given to *Irak unie* since the final *-e* of the adjective *unie* (*united*) is inappropriate for the masculine *Irak*.

EXPERIMENT 3

Using the gender verification task, we can see whether the vowel/consonant difference that was observed in Experiment 2 is maintained even when the gender is not explicitly required. If the two tasks are equally tapping into the mechanisms involved in generating the gender of a noun, then it should be more difficult to verify the form of an adjective that follows a place-name beginning with a vowel than one beginning with a consonant. The difficulty in determining that *Israël* is masculine putatively arises from the fact that the article that accompanies it provides no information about gender (and that the orthographic cues to gender are not given much weight) and this is true regardless of the task.

On the other hand, if one can draw upon knowledge about the form of adjectives that have been encountered in association with the place-name, then gender verification might be unaffected by vowel onset because the response could be made on the basis of specific associations set up between the place-name and the presented adjective or an analogous one. For example, while an association might never have been set up between *Australie* and *profonde* because they had never been encountered together before. It is quite possible that one might have been set up between *Australie* and *lointaine* (*far*) from which the grammatical appropriateness of putting *profonde* with *Australie* can be determined by the analogous use of the feminine form.

What is additionally interesting about the gender verification task is that we can look at a situation where the article becomes potentially inform-

ative even when the place-name begins with a vowel. In French, most adjectives are placed after the noun that they qualify, but there are a few which are standardly placed before. If this latter type of adjective begins with a consonant, then the definite article will be expanded to its full form, hence revealing its gender. For example, the adjective *joli* (*pretty*) comes before a noun and therefore we get *le joli Israël* where the masculine article *le* manifests itself. This means that if gender information at least partly arises from knowledge of the article, as suggested by the vowel/consonant difference observed in Experiment 2, then we might expect a reduction in this difference when the adjective is one that precedes the noun because, here, any article that is generated must be informative of gender.

So in Experiment 3, there were two conditions for each place-name. The name was presented either with a preceding adjective (e.g., *joli Israël, haute Suisse*) or with a following adjective (e.g., *Israël froid, Suisse forte*). For each place-name beginning with a vowel, there was a matched place-name beginning with a consonant (e.g., *joli Israël/joli Brésil, Israël froid/Brésil froid*). The question was whether the vowel/consonant difference that was observed in the gender decision task of Experiment 2 would reveal itself in a gender verification task and whether this would be more so when the adjective followed the noun than when it preceded it.

Method

Subjects

Thirty native French speakers from the same subject pool as the other experiments participated in the experiment. None had been a subject in the previous experiments.

Materials and Procedure

The same pairs of place-names that were used in Experiment 2 were used here, but presented in combination with an adjective. As described in Experiment 2, the pairs of place-names were matched on rated familiarity and orthographic ending, but differed on whether they began with a vowel or a consonant (e.g., *Israël* and *Brésil*).

Half of the adjectives that were used were ones which come before the noun that they modify and were presented as such (e.g., *joli Israël, nouvelle Thaïlande*), while the other half were ones which come after the noun (e.g., *Israël froid, Thaïlande profonde*). Owing to a shortage of adjectives which were appropriate for describing place-names and which had different forms for the two genders, each adjective was used more than once. Each member of a vowel/consonant pair was combined with the same adjective (e.g., *joli*

Israël/joli Brésil; *Israël froid/Brésil froid*). An attempt was made to use an adjective which was just as semantically plausible for each member of a pair (e.g., neither Israel or Brazil are typically associated with being "cold," and both might be considered equally "pretty").

The items were split between two groups of 15 subjects such that any one subject saw each place-name only once. For example, one group saw *joli Israël* and *Brésil froid*, while the other group saw *joli Brésil* and *Israël froid*. Each group received the same 36 distractor items which required a *no* response. These consisted of place-names that were not used for the *yes* responses (e.g., *Ukraine profond*, *jolie Colorado*). No articles were used.

Items were presented to subjects following the same procedure as in the previous experiments. The adjective and noun were placed side by side in the order appropriate for that adjective. Subjects were instructed to decide by button-press whether the adjective and noun went together grammatically. It was emphasized that the response should be based on whether the agreement between the adjective and noun was grammatically correct and not on whether the subject thought that the adjective was an appropriate description of the place it was combined with. Subjects were asked to respond as quickly but as accurately as possible.

Results

Table III presents the gender verification times and error rates. In this experiment, masculine and feminine gender were collapsed together because, otherwise, the reaction time (RT) for each condition would be based on a mean of only five items at the maximum.

The main effect of type of onset (vowel vs. consonant) proved to be significant both for reaction time, $F_1(1, 29) = 9.71$, $p < .01$; $F_2(1, 17) = 11.38$, $p < .01$, and error rate, $F_1(1, 29) = 6.79$, $p < .02$; $F_2(1, 17) = 4.11$, $p < .1$. When the adjective came before the noun, error rates were higher

Table III. Gender Verification Times (Reaction Time, or RT, in Milliseconds) and Percentage Error Rates (ER) for Experiment 3; Standard Error is in Parentheses

	Vowel onset			Consonant onset		
	Example	RT	ER	Example	RT	ER
Adjective after	Israël froid	1381 (58.8)	7.4% (.27)	Brésil froid	1240 (59.1)	4.8% (.23)
Adjective before	joli Israël	1381 (78.6)	13.7% (.52)	joli Brésil	1291 (47.4)	7.8% (.36)

than when the adjective came after the noun, $F_1(1, 29) = 6.28$, $p < .02$; $F_2(1, 17) = 5.26$, $p < .05$, but response times did not differ, Fs < 1.

There was no interaction between the position of the adjective and onset type, though the subject analysis of response times was of marginal significance, $F_1(1, 29) = 3.84$; $p < .1$; $F_2(1, 17) = 0.24$, $p > .1$ for RT and $F_1(1, 29) = 1.44$, $p >, .1$; $F_2(1, 17) = 1.22$, $p > .1$ for errors. Note that any tendency for an interaction on error rate was in the direction of there being a larger vowel/consonant effect when the adjective came before the noun rather than after, while the reverse tendency was seen in reaction time. Taking into consideration the complete lack of an interaction on the item analysis of RTs it seems safe to say that the onset effect held as much for adjective-noun pairings as it did for noun-adjective pairings.

Discussion

The same effect of onset was found in the gender verification task as was found in the gender decision task. Therefore, it is clear that a vowel onset makes the gender of a place-name difficult to determine and this is presumably because it prevents information about gender being directly generated from the article that goes with the word. Surprisingly, however, the vowel/consonant difference was maintained when the adjective came before the noun even though this means that the article now revealed the gender of the noun. The implication of this is that the only form of the article that was consulted for the determination of gender was that directly attached to the noun. That is, it was the *l'* of *l'Israël* that was considered when trying to determine gender rather than the *le* which emerged when an adjective intervened between the article and the noun. It seems that when the full form of the article is required (i.e., before an adjective beginning with a consonant), it is generated from the same information that is used to determine the gender of the word, rather than being explicitly available in lexical memory as an underlying representation of *l'*.

More errors were made when the adjective came before the place-name than after it regardless of whether the place-name began with a vowel or consonant. One possible explanation for this is that analysis of the gender of the phrase centers upon the gender of the noun rather than the adjective because the adjective can potentially be of either gender while the gender of the noun is unique. So, when the adjective is read before the noun, the gender of the adjective is sometimes lost during the pursuit of the gender of the noun, whereas if the adjective is processed after the noun, the gender that has been extracted from the noun merely needs to be confirmed against the form of the adjective. Alternatively, there may be a more a simplistic perceptual explanation in that the ending of the adjective, which is where

the only indication of its gender can be found, is easier to perceive when it occurs at the end of the phrase rather than in the middle.

GENERAL DISCUSSION

The three experiments reported here looked at several possible sources of gender information in French. The nature of the word-form, in terms of the word-ending, was shown in Experiment 1 to be involved in gender determination, as was also shown by Bates et al. (1995), Desrochers et al. (1989), and Tucker et al. (1977). Experiments 2 and 3, however, demonstrated that there was some difficulty in using word-endings to determine gender in a situation where it was the most predictive source of information available. These experiments further revealed that knowledge of the article that can go with the noun provided information that played a role in gender determination. When the article was uninformative of gender, as in the case of place-names starting with a vowel, responses became more difficult. Finally, the influence of word-form and article informativeness were the same for high-frequency and low-frequency words.

What is the best way to account for this pattern of results? From the results of Experiment 2, there seem to be two candidate explanations. The first is the ''dual-source'' explanation whereby gender assignment is based on information associated with the lexical representation of the noun (e.g., the articles that can go with it), but is also influenced by nonlexical information based on word-form (e.g., *-ie* means ''feminine''; *-t* means ''masculine''). If this influence only arises after the lexical information is extracted (perhaps in a final checking stage), regularity and consistency effects can be explained, as well as their failure to interact with frequency. Furthermore, if it is difficult to extract gender from lexical information, as in the case of place-names beginning with a vowel, then the use of nonlexical information must wait. Therefore, high predictiveness of orthographic endings does not compensate for the difficulty of using lexical information and this leads to the observed difficulty in determining the gender of place-names beginning with a vowel.

The problem with the dual-source model, however, is that it does not explain why the orthographic information would only be used after lexical information is extracted. What function does it serve? For words which are associated with at least one article that provides information about gender (and this is virtually all words except place-names beginning with a vowel), the article is a 100% reliable guide to gender. There is no need for the orthographic information to confirm the gender, and in fact it is potentially a distraction given that it might conflict with the information extracted from

the article (as with irregular and nonpredictive words). Therefore, while the dual-source model is able to explain the pattern of data that has been obtained in grammatical gender research, it requires assumptions that are motivated solely by the data themselves.

The second explanation does not require such assumptions. The neural network approach provides an account of the regularity effect in a quasiregular domain as well as the vowel/consonant difference with place-names, as long the output units represent articles and other associated information rather than abstract gender labels. For example, during training (i.e., during exposure to the language) an association between the noun *folie* and the articles *une* and *la* is learnt, so that when *folie* is the input, the articles *une* and *la* become activated as output from which the feminine gender can be derived with 100% certainty. If the article does not provide any information about gender (as in the case of place-names beginning with a vowel), some other strategy must be used—for example, drawing upon links that have been formed as a result of exposure to the place-name in a context which reveals the gender (e.g., *elle est belle*, *l'Australie*). In this way, the results of Experiment 2 can be explained.

The regularity and consistency effects arise from the connections between the orthographic input units and the article output units. That is, the system is influenced by the fact that final *-ie* is typically associated with the occurrence of the articles *une* and *la*, that final *-t* is typically associated with the occurrence of the articles *un* and *le*, and that final *-le* following a consonant is associated mostly with *un* and *le* but sometimes with *une* and *la*. Thus, as with any quasiregular domain, the generation of output will be slowed by irregularity and inconsistency.

This explanation crucially differs from the dual-source account in terms of the locus of the effect of orthographic factors. In the neural network account, orthography has its effects in the course of activating information about the article, whereas in the dual-source account, orthographic factors are independent of article information and come into play only after such lexical information has been activated. The dual-source account can therefore easily handle the lack of interaction between word frequency and regularity (and consistency), while this is a problem for the neural network account. It should be pointed out, though, that Desrochers et al. (1989) did find a larger effect of predictiveness for low-frequency words than for high-frequency words in relation to error rates and this is difficult for the dual-source model to explain.

So the neural network approach seems to provide a more natural way to explain the regularity and vowel/consonant effects than does the dual-source account, but it does not readily explain the lack of interaction of regularity with word frequency. Undoubtedly a neural network could be

designed so that no interaction between frequency and orthographic factors would occur, but it needs to explain at the same time why an interaction does occur when regularity is defined in terms of grapheme-phoneme relationships. The answer might lie in statistical differences in the relationship between orthographic input and phonological output and the relationship between orthographic input and article output, perhaps related to the fact that the latter involves only two types of output while the former involves many.

Interestingly, there is some evidence that even grapheme-phoneme regularity does not interact with word frequency in French (Content, 1991). So it is possible that there is something special about the statistical facts of French in this regard. One possibility (see Content) is that high-frequency words are not as high in frequency in French as they are in English, though it is not obvious why this should be so. The mean high-frequency value in Experiment 1 was 103 per million, while the mean high-frequency value in the study by Seidenberg et al. (1984) was not much more (137 per million), yet regularity interacted with frequency. This issue remains to be resolved.

In sum, then, the experiments reported here showed that grammatical gender is not something that French speakers mentally store together with each word in an abstract form. Instead, representation of the article appears to play a role in determining gender, but so too does the orthographic information contained at the end of the word. A neural network model might turn out to be the best way to capture these relationships and, as such, gender can be seen as a useful quasiregular domain in which to explore that approach.

APPENDIX

The following are the pairs of items used in Experiment 1, regular followed by irregular.

		Low frequency
Masculine:	socialisme	squelette
	maquereau	magazine
	rugby	renne
	four	foie
	guidon	gorille
	museau	mercure
	paludisme	dividende
	plongeoir	pamplemousse
	carnage	carrosse
	bambou	bastion

Feminine:	palette	perdrix
	toupie	tumeur
	salade	souris
	ficelle	fourmi
	touffe	tribu
	taie	glu
	dose	dot
	cascade	cuiller
	romance	*rançon*
	fixité	frayeur

High frequency

Masculine:	minuit	malaise
	détail	domaine
	incident	incendie
	genou	génie
	travail	silence
	jour	monde
	argument	intervalle
	message	murmure
	concret	comité
	ouvrage	arrière
Feminine:	envie	erreur
	pluie	peau
	soie	soif
	action	image
	figure	vertu
	folie	forêt
	peinture	plupart
	famille	façon
	vie	fois
	classe	cour

The following are the pairs of place-names used in Experiments 2 and 3, vowel onset followed by consonant onset. The adjectives used in Experiment 3 are presented for each pair of place-names, adjectives that come after the noun followed by adjectives that come before the noun.

Masculine:	Oklahoma	Montana	vert	joli
	Ouganda	Guatemala	lointain	petit
	Irak	Danemark	uni	petit
	Israël	Brésil	froid	joli
	Afghanistan	Pakistan	lointain	grand

	Oregon	Michigan	vert	petit
	Arkansas	Kansas	sec	petit
	Uruguay	Paraguay	fort	joli
	Angola	Nebraska	uni	haut
Feminine:	Allemagne	Bretagne	froide	vieille
	Islande	Thaïlande	profonde	nouvelle
	Auvergne	Pologne	profonde	belle
	Aquitaine	Lorraine	chaude	vieille
	Argentine	Chine	lointaine	nouvelle
	Écosse	Suisse	forte	haute
	Estonie	Lituanie	froide	petite
	Algérie	Tunisie	chaude	jolie
	Australie	Russie	profonde	belle

REFERENCES

Andrews, S. (1982). Phonological recoding: Is the regularity effect consistent? *Memory & Cognition*, *10*, 565–575.

Bates, E., Devescovi, A., Pizzamiglio, L., D'Amico, S., & Hernadez, A. (1995). Gender and lexical access in Italian. *Perception & Psychophysics*, *57*, 847–862.

Content, A. (1991). The effect of spelling-to-sound regularity on naming in French. *Psychological Research*, *53*, 3–12.

Desrochers, A., Paivio, A., & Desrochers, S. (1989). L'effet de la fréquence d'usage des noms inanimés et de la valeur prédictive de leur terminaison sur l'identification du genre grammatical. *Revue Canadienne de Psychologie*, *43*, 62–73.

Deutsch, W., & Wijnen, F. (1985). The article's noun and the noun's article: Exploration into the representation and access of linguistic gender in Dutch. *Linguistics*, *23*, 793–810.

Glushko, R. J. (1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 674–691.

MacWhinney, B., Leinbach, J., Taraban, R., & McDonald, J. (1989). Language learning: Cues or rule? *Journal of Memory and Language*, *28*, 255–277.

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.

Radeau, M., Mousty, P., & Bertelson, P. (1989). The effect of the uniqueness point in spoken-word recognition. *Psychological Research*, *51*, 123–128.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.

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.

Stanhope N., & Parkin, A. J. (1987). Further explorations of the consistency effect in word and nonword pronunciation. *Memory & Cognition*, *15*, 169–179.

Trésor de la langue française: Dictionnaire des fréquencies. (1971). Paris: Klinksieck.

Tucker, G. R., Lambert, W. E., & Rigault, A. A. (1977). *The French speaker's skill with grammatical gender: An example of rule-governed behavior*. The Hague: Mouton.

Tucker, G. R., Lambert, W. E., Rigault, A. A., & Segalowitz, N. (1968). A psychological investigation of French speaker's skill with grammatical gender. *Journal of Verbal Learning and Verbal Behavior*, *7*, 312–316.