

DEVELOPMENT OF GENDER CLASSIFICATIONS: MODELING THE
HISTORICAL CHANGE FROM LATIN TO FRENCH

MARIA POLINSKY

EZRA VAN EVERBROECK

University of California, San Diego

University of California, San Diego

We present and analyze the results of a connectionist simulation which modeled the reanalysis of the Latin gender system in its transition to Old French. The network reanalysis was based solely on formal cues (word endings and analogy with other words) and on frequency. The results are in accordance with the historical data, and certain errors in simulations are also amenable to principled explanations. Simulations improve dramatically when the networks incorporate information about the Celtic substrate which presumably interfered with gender assignment in Gallo-Romance. This finding has a bearing on issues of gender assignment and processing in bilinguals. Simulations also improve with the introduction of more elaborate recurrent networks, which suggests implications for future connectionist modeling. In particular, the results could be applied to the modeling of gender change in other Romance languages and to the modeling of comparative Romance gender systems. The method proposed here would be advantageous for such simulations since it allows the modeler to take into account a rich variety of facts reflecting actual linguistic history.*

1. INTRODUCTION. 'Genders are classes of nouns reflected in the behavior of associated words' (Hockett 1958:231).¹ How a language divides its noun lexicon into genders (noun classes) is a complicated question, centering on the interaction between formal and semantic bases of noun classification. A large body of work has been devoted to the semantic or pragmatic base of gender, and here the accounts range from very complicated, culture-based systems proposed for some Australian languages (Lakoff 1987, Mylne 1995, Walsh 1993, Harvey & Reid 1997) to the demonstration of the relevance of superordinate semantic categories in gender assignment (Mandler 2000). On the other hand, gender assignment has also been accounted for by a combination of formal cues (see Corbett 1991 for a useful overview), among which morphophonemic cues and frequency are most important. Of course, articulated formal systems of gender classification include at least one robust semantic feature, for instance, natural gender as in many Indo-European languages, or animacy and the distinction between living and nonliving entities as in Algonquian languages (Joseph 1979, Lapointe 1988:83, Corbett 1991). The appeal to formal cues and frequency is particularly strong given that gender is acquired fairly early, and young language learners are sensitive to the outward appearance and statistical distribution of linguistic elements (Jusczyk et al. 1993, 1994, Saffran, Aslin & Newport 1996, Saffran, Newport & Aslin 1996). In the acquisition literature, it is now the predominant theoretical view that children pay much more attention to intralinguistic information than to semantic information in determining gender (Demuth et al. 1986, Karmiloff-Smith 1979, Berman 1985). Similar findings,

* This work was supported in part by a grant from the University of California Academic Senate. We are indebted to Nicoleta Bateman, Joan Bybee, Bernard Comrie, Greville Corbett, Gary Cottrell, Jeff Elman, Mary Hare, Dan Jackson, Dan Jurafsky, Alex Yanovsky, two anonymous referees, and our associate editor Donna Gerdtz for helpful comments on this paper. We are particularly grateful to Orin Gensler and Brian Joseph, who went over the paper many times and have given us numerous suggestions on how to improve it.

¹ The quote from Hockett indicates that 'gender' and 'noun class' are two terms denoting the same concept; the choice of the term depends on a particular terminological tradition associated with a given language or language family (see Corbett 1991:3-5 for a discussion of terminology). In this paper, we use both terms interchangeably.

although in a weaker form, pertain to the acquisition of gender by adult L2 learners (e.g. Andersen 1984 on L2 learning in general and Sokolik & Smith 1992 on L2 French).

In this paper, we investigate THE EXTENT to which the restructuring of noun class (gender) systems can be explained using the type of morphophonemic and frequency information that is accessible in acquisition and language contact. We argue that in the cases which we analyze, the restructuring of noun classes can be accounted for on the basis of the morphophonemic similarities of items as well as their type and token frequency.² This is exactly the type of information to which connectionist simulations are known to be sensitive (Plunkett & Marchman 1991, Hare & Elman 1995). The case we examine is the restructuring of the gender system in Romance. In particular, we show how a sophisticated connectionist network was applied to model the change from the three-gender system found in Latin to the two-gender system of Old French. The importance of this project is twofold: first, it is a good tool for testing hypotheses concerning the interaction of formal and semantic cues in gender assignment in natural language; second, it provides us with new connectionist insights, in particular, with ways of implementing more complex simulations which are closer to actual language situations.

1.1. A GENERAL PERSPECTIVE ON GENDER. Our major preoccupation is with the change and restructuring of gender in a natural language. To address this question, we first examine the general mechanisms of gender assignment. We are not concerned here with the ways in which gender may arise in natural language—this is a question for a very different line of research and we do not believe the simulations presented below can be of any use in addressing this issue.

The linguistic complexity of gender comes from three different sources: the number of genders within a given language, the nature of gender agreement, and the motivation for assigning gender to a particular word. Gender systems range in number from the well-known two- or three-gender systems of Indo-European languages to the particularly rich Niger-Congo systems (with as many as twenty-two genders in some languages, depending on the criteria used to identify categories). A large body of theoretical work deals with the larger, more ‘exotic’ systems (see Corbett 1991 for an overview), but experimental studies have largely concentrated on the smaller systems of Indo-European languages (see van Berkum et al. 1999 for a useful overview).

In all languages that have gender, its role within the grammar is central, controlling various forms of agreement. Recall Hockett’s definition above: gender is an inherent feature of a noun which is reflected in the form assumed by one or more constituents, usually within the same clause. Prenominal modifiers showing gender agreement can prime the nouns they modify and thus can facilitate lexical access (Wicha 2002, Wicha et al. 2003). From the point of view of the speaker, gender provides support for the construction of utterances. From the point of view of the hearer, gender agreement can provide important cues to word identity, coreference, and syntactic structure. Such cues may include the phonological shape of agreement markers, morphological rules determining the placement of agreement elements, and the very order of agreeing elements in relation to the noun with which they agree.

² Our interest in such information does not imply that we assume that only children can innovate and initiate language change. Adult usage is often responsible for language change and reanalysis; evidence of this comes from second language acquisition, creole studies (e.g. Sankoff 1980 on relativizers in Tok Pisin), cultural categorization (Zubin & Köpcke 1986), and substrate and superstrate changes (Thomason & Kaufman 1988, among many others; see also §2.4).

In assigning gender to a particular noun, semantic and formal factors can both play a role. Some of the best-known instances of semantically motivated gender assignment are found in certain Dravidian languages, where all nouns describing males are masculine, nouns describing females are feminine, and everything else is neuter (Corbett 1991:7–12). Russian or French can be cited as good examples of formally motivated gender assignment, but even those gender systems have a partial semantic base (see Corbett 1991:57ff. and Lapointe 1988 for a general overview, Corbett 1982 for Russian gender, Bidot 1925, Mel'čuk 1974, Tucker et al. 1977 for French gender). In fact, most languages that have gender are mixed systems in which gender assignment is based partially on semantics and partially on formal characteristics.

The semantic characteristics involved in gender assignment are typically of a straightforward, general nature (human/nonhuman, male/female, moving/nonmoving, edible/inedible). Such generality is understandable if one bears in mind that the primary learners of gender systems are young children who cannot yet have access to complex semantic and cultural information of the sort to which some accounts of gender systems have appealed (see Polinsky & Jackson 1999 for a discussion). The formal characteristics involved in gender assignment can be classified, roughly speaking, into three major types: suprasegmental phonological (e.g. Afar, per Corbett 1991:51f.), segmental phonological (e.g. Yimas, French, Godie — see Corbett 1991:53–61), and morphological (e.g. Russian). The placement of segmental cues constitutes another variable of gender assignment: cues can occur at the onset of the word, as in some languages of the Caucasus (Nichols 1989), or at the end of the word, as in Yimas or many Indo-European languages (Smoczyńska 1985, Zubin & Köpcke 1986, among others).

Much discussion of grammatical gender has been directed at the psychological reality of gender assignment: are the putative semantic and formal principles that motivate gender assignment in any particular language alive in the mind of a speaker? Even if the answer to this question is yes (as suggested by many researchers), some motivations are part of the history of that language and may be currently inaccessible to language users. If a certain principle of gender assignment is obscured by language change, this may result in opacity. Opacity in gender assignment in turn leads to the need for rote memorization of the gender of certain nouns. Of course there is always a closed set of nouns (usually short and frequent ones) that need to be learned by rote, but if the number of such words pushes towards a critical mass, they may undergo a reanalysis, assimilating to those nouns whose gender is clear and/or they may foster the reanalysis of the entire gender system (including the loss of gender distinctions altogether).³ Thus at any given time, a gender system is subject to the pressures of change and to the pressures of analogy, which do not always produce the same result. The tendency to eliminate irregularity can be viewed as a major driving force in the restructuring of gender systems: the gender system changes in such a manner that new cues become

³ At least two separate issues arise with respect to this line of reasoning. The first is the question of what constitutes the critical mass and whether type or token frequencies have a direct effect on the need to assimilate nouns or restructure the system. This issue is beyond the scope of this article, but we touch upon it below in our discussion of the differences in the contribution made by high- and mid-frequency nouns. The second issue concerns predicting the outcome of increased opacity: if a number of nouns must be learned by rote, what if anything determines the assimilation of these nouns versus the general restructuring of the system? For instance, in Romance, the transition from Latin to French, Italian, and Spanish involved a complete restructuring (three > two genders), whereas in Rumanian, neuters were retained, which effectively led to a clear differentiation of gender types in the singular vs. the plural. Such a bifurcation suggests that a definitive answer to this question may be unattainable, but definitely worth pursuing.

available to a speaker and provide a more transparent rationale for categorization, thus facilitating gender assignment.

The transition we analyze is the change from the three-gender system of Latin to the two-gender system of Old French. We believe that the Latin system became increasingly complicated, to the point that reanalysis into the simpler system of Old French was easier than learning and maintaining the old system. We discuss the specifics of both gender systems in the next section. Here, we would like to offer some general remarks on the crucial differences between a three- and a two-gender system. There are at least three reasons why the reduction of genders from three to two leads to a SIGNIFICANTLY different gender assignment system. First, two-gender systems can treat gender assignment as a binary decision, with one gender serving as the default form, for example, for new words entering the language. Because three-gender systems are not binary, the assignment of gender is a more complex problem as it may involve more steps in the computation of gender. Second, many two-gender systems have a roughly equal balance in both type and token frequency across the two genders (although type frequencies can be greater for one gender). By contrast, there are often large disparities in type and token frequency within three-gender systems, with much lower type frequencies in one gender. As we discuss below, the neuter gender in Latin was part of such a system. Third, the existence of three or more genders can lead to complications in the formal marking on the noun itself and on agreeing elements; over time, some agreement markers and some gender assignment cues may fail to distinguish genders due to, for example, phonological change and subsequent reanalysis. This latter possibility played a crucial role in the restructuring of Romance gender. We now turn to the specifics of this process.

2. FROM LATE LATIN TO OLD FRENCH. In this section we provide a simplified overview of the gender system in Latin and French and of the changes that led to the development of the Old French gender system. The reader interested in a detailed history of French should refer to a large body of philological literature, some of which is cited here.

2.1. BASIC HISTORICAL INFORMATION. French belongs to the Romance group of the Indo-European family; it developed from the Latin of Transalpine Gaul, which showed early signs of dialectal differentiation into a northern and central variety on the one hand, and a southern variety on the other. The former dialect group is known as *langue d'oïl*, the latter as *langue d'oc* (Occitan). Standard Modern French is based largely on *langue d'oïl*, and it is this variety that we examine in this paper. Although *langue d'oïl* and *langue d'oc* did not differ crucially in their gender systems, they had quite different sound systems and grammars. Old French and Old Occitan were closer to each other than their modern descendants, but attempting to model change across both dialects at once would be unrealistic.

The origins of the language known now as French begin with the Roman colonization of Gaul in the second half of the second century BC. At this point the Romans occupied the southern part of modern France. During Caesar's time (mid-first century BC), Gaul lost its independence, and Roman soldiers—the majority of whom were speakers of colloquial Latin—occupied most of the territory. By the reign of Claudius (14–37 AD), Gaulish senators were already represented in the Roman Senate. Latin started its aggressive spread into the linguistic milieu of Gaul, competing with Gaulish, the Continental Celtic language spoken by the indigenous population.

At the end of the fifth century, a Catholic Frankish kingdom under Clovis was established in northern Gaul, with the capital in Tours. Under Clovis, Latin came to

be used as the language of religion and administration. The majority of the population spoke a Romance vernacular with a substantial Celtic substrate, and Clovis's reign is the conventional watershed at which the precursor to Old French gains the upper hand over Celtic. The first written document in Early Old French or Gallo-Romance appeared in 842 (*Les serments de Strasbourg* 'the Strassburg Oaths'), followed by short religious poems ca. 880 (*Poème sur sainte Eulalie*; *Passion du Christ*; *Vie de saint Léger*).⁴ Overall, the period between 850 and 1050, when few literary texts appear, is considered Early Old French. Old French proper spans the eleventh through thirteenth centuries, culminating in such well-known literary documents as *Chanson de Roland* (end of the eleventh century) and *Roman de la Rose* (end of the thirteenth century). These and other texts provide us with a very clear and detailed picture of the Old French grammar and lexicon.⁵

2.2. DECLENSION AND GENDER IN LATIN AND OLD FRENCH. CLASSICAL LATIN has three genders: masculine, feminine, and neuter.⁶ Latin nouns trigger agreement in gender on demonstratives, adjectives, participles, and numerals one through three. For example:⁷

- (1) Gender agreement in Latin
 - a. ille bonus amiicus
that-M good-M friend-M.NOM.SG
 - b. illa bona silva
that-F good-F forest-F.NOM.SG
 - c. illud bonum tempus
that-N good-N time-N.NOM.SG
 - d. tres laudatii ooraatoorees
three-M praised-M speaker-M.NOM.PL
 - e. tres laudatae viae
three-F praised-F road-F.NOM.PL
 - f. tria laudata bella
three-N praised-N war-N.NOM.PL

Turning now to gender assignment principles, Latin masculine and feminine nouns have a semantic core, which is reflected in the SEMANTIC ASSIGNMENT RULES:

- A. nouns denoting male humans and supernatural beings are masculine, regardless of their phonological shape, e.g. *scriba* 'scribe', *filius* 'son', *rex* 'king';

⁴ The surviving manuscript of the Oaths actually dates from 150 years later, which may have affected the language recorded in it.

⁵ For a detailed history of French, see Meyer-Lübke 1923, Cohen 1973 and further references there on pp. 427–35, Wartburg 1888, Pope 1961, Elcock 1960, Wright 1982, Rohlfs 1960, among many others. A brief overview of gender changes throughout the history of French can be found in Harna 2000.

⁶ The relationship between Old Latin, Classical Latin, and Vulgar Latin is not without controversy. Some researchers believe that Classical Latin was an ephemeral standard that, although associated with a vast literature, was not actually widely used (Meillet 1966, Baldi 1999). On this view, the transition to Vulgar Latin was from Old, not Classical, Latin. This complex relationship among the three arguably different varieties of Latin does not directly affect the general design of our study, and we do not concentrate on it below unless specifically warranted by individual lexical items.

⁷ Here and below, the spelling reflects actual pronunciation, not traditional orthography; long vowels are represented by doubling.

- B. nouns denoting female humans and supernatural beings are feminine, regardless of their phonological shape, e.g. *dea* ‘goddess’, *filia* ‘daughter’, *soror* ‘sister’.

Most of the nouns, which obviously do not denote male or female humans, are distributed over the three genders based on their declension class, and thus on morphological factors.

Nouns inflect for case, with five main cases: nominative, accusative, genitive, dative, and ablative; one could also recognize the vocative and maybe the locative, but since these two cases are not relevant for the discussion below, we do not include them. Traditional grammars distinguish between the five declension classes illustrated in Table 1. Classes II and III have more subtypes, but we ignore them for simplicity. The table is also simplified in that it does not include neuter nouns.

	I (<i>a</i> -stems)	II (<i>o</i> -stems)	IIIa (consonantal stems)	IIIb (<i>i</i> -stems)	IV (<i>u</i> -stems)	V (<i>e</i> -stems)
	capra- ‘goat’	amiico- ‘friend’	ooraatoor- ‘orator’	urbi- ‘city’	fructu- ‘fruit’	die- ‘day’
SINGULAR:						
Nominative	capra	amiicus	ooraator	urbs	fructus	dies
Genitive	caprai	amiicii	ooraatooris	urbis	fructuus	diei
Dative	caprai	amiicoo	ooraatoorii	urbii	fructuu/ fructuui	diei
Accusative	capram	amiicum	ooraatoorem	urbem	fructum	diem
Ablative	capraa	amiicoo	ooraatoore	urbe	fructuu	diee
PLURAL:						
Nominative	caprai	amiicii	ooraatoorees	urbees	fructuus	dies
Genitive	capraarum	amiicoorum	ooraatoorum	urbium	fructuum	dierum
Dative	capriis	amiiciis	ooraatooribus	urbibus	fructibus	diebus
Accusative	capraas	amiicoos	ooraatoorees	urbiis/ urbees	fructuus	dies
Ablative	capriis	amiiciis	ooraatooribus	urbibus	fructibus	diebus

TABLE 1. Declension classes in Classical Latin.

Even a cursory look at the declension forms reveals a significant homophony of case forms within and across the declension classes. In the singular, the genitive and the dative have the same form in classes I and V; the dative and the ablative have the same form in class IV. In the plural, the dative and ablative are identical for all classes; the nominative and accusative coincide for classes III–V.⁸ Declension classes I and II have the nominative plural in *-i*; all the other classes have a nominative plural which ends in *-s*.

If we now bring neuter nouns into the picture, these nouns never distinguish between the nominative and the accusative, and their plural nominative/accusative always ends in *-a*, regardless of the declension class; for example: *bellum*, *bella* ‘war’ (II); *tempus*, *tempora* ‘time’ (III); *cornuu*, *cornua* ‘horn’ (IV). Despite the homophony of forms, the declension class of most nouns can be determined on the basis of the nominative singular; class III nouns also require access to the genitive singular, and this case form is needed to distinguish between class IV and class II nouns in *-us*.

⁸ The nominative/accusative syncretism in the plural is actually greater, because neuter nouns have identical nominative and accusative plural regardless of the declension.

The gender of a given noun correlates with its declension class, but this morphological correlation holds only for nouns that are not subject to semantic assignment rules. Whereas the large declension class III contains nouns of all genders, the following two CORRELATIONS apply to the other classes:

- A. nouns of declension classes I and V are predominantly feminine;
- B. nouns of declension classes II and IV are predominantly masculine or neuter.

Both correlations are subject to exceptions. Such exceptions are usually given as lists in Latin grammars; the implicit assumption is that they have to be learned by rote. The gender of 'exceptional' nouns was unstable and was subject to hesitation already in Old Latin (Ernout 1927, 1953, Pinkster 1990). By the Late Latin period, when speakers of the language had become more numerous, the number of bilingual speakers had grown significantly, and a socially validated coherent standard had become much less evident, there was even more hesitation over the gender assignment of individual nouns.

The size of gender and declension classes in Latin varied significantly. Starting with declensions, classes IV and V were quite small. Class V nouns started acquiring parallel class I declension forms already in Old/Classical Latin; likewise, class IV nouns started acquiring parallel class II declension forms. Class III was the largest declension class. Table 2 presents the statistics on declension class size.⁹ With respect to genders, the neuter is the smallest class; the overall statistics are given in Table 3.

CLASS I	CLASS II	CLASS III	CLASS IV	CLASS V
21.6%	23.7%	52.6%	1.4%	0.7%

TABLE 2. Declension class size as a percentage of the nominal lexicon in Classical Latin (1,500 lexical items).

MASCULINE	FEMININE	NEUTER
38.1%	40.8%	21.1%

TABLE 3. Gender class size in Classical Latin (1,500 lexical items).

All this is a presentation of a largely idealized system, sustained by classical conservatism but not entirely supported by the colloquial language. The variation in stem types starts early on; nouns whose declension did not quite fit the standard templates were subject to inflectional variation, which reveals general hesitation over correct forms. Evidence for the increasing disparity between the literary and the spoken language comes from the language of inscriptions which, from Caesar's time on, shows a large number of deviations from the standard (Ernout 1927, Vaananen 1981), and from the regional variation that became more apparent as the Roman Empire grew.

Specifically, Late Latin is characterized by the following changes:

- A. the case system undergoes a significant reduction due to case syncretism and the increased use of prepositional phrases in lieu of synthetic case forms; i.e. the dative and ablative cases disappear;

⁹ Here and below, the statistics for Classical Latin are averages based on two sources: frequency data in Diederich 1939 and a random sampling of 1,500 nouns from Classical Latin texts (PHI CD ROM #5.3, Los Altos, CA: Packard Humanities Institute, 1991). At the beginning of our project, in 1997, the Perseus Project data on Latin words and word counts were not yet available.

- B. the number of declension classes becomes smaller; in particular, the small classes (IV, V) become even smaller, and most words of these classes shift into classes II and I, respectively;
- C. the neuter gender decreases from 21.1% to 18.4% of the nominal vocabulary, and many neuters shift into the feminine gender, which grows to 43.1%;
- D. thus, the final development of nominal forms in Late Latin can be represented as in Table 4 (compare with Table 1). We discuss the relevant phonological changes below.¹⁰

	I	II	III
SINGULAR:			
Nominative	capra	amicu(s)	urb(s)
Genitive	caprae [ai]	amici	urbi
Accusative	capra(m)	amicu(m)	urbe(m)
PLURAL:			
Nominative	caprae [ai]	amici	urbe(s)
(Genitive)	capra(rum)	amico(rum)	urbi(um)
Accusative	capra(s)	amico(s)	urbe(s)

TABLE 4. Declension classes in Late Latin (the vocative case is not shown). The elements in parentheses were deleted in pronunciation from the early stages of Late Latin. Loss of vowel length is also reflected here.

In Early Old French, the nominal system is further restructured. The genitive becomes less robust as a category and eventually disappears. Further leveling of paradigms reduces the number of declension classes. The system that emerges has two cases (nominative and accusative) and two genders (masculine and feminine).

2.3. SUMMARY OF MAJOR CHANGES IN THE NOMINAL PARADIGM FROM LATIN TO OLD FRENCH. Simplifying to a considerable degree, the main changes in the nominal paradigm can be subdivided into systematic changes, changes that involved a subclass of nouns, and changes that affected individual lexical items. The major systematic changes that affected gender assignment included:

- A. Declension mergers:
- (i) Declension class II absorbs declension class IV (Late Latin);
 - (ii) Declension class I absorbs declension class V (Late Latin);
 - (iii) Declension class III plurals assimilate to declension class II plurals (Old French).
- B. Case loss:
- (i) Latin Nom., Acc., Gen., Dat., Abl. > Old French Nom., Acc., (Gen.).
- C. Gender reanalyses:
- (i) Declension III neuters in *-us* are absorbed by declension II nouns, most becoming masculine, as in example 2 below;
 - (ii) Neuter singulars are reanalyzed as masculine singular, as in 3;
 - (iii) Neuter plurals are reanalyzed as feminine singular, as in 4.

¹⁰ The statistics for Late Latin are based on St. Jerome's Vulgate: the Old Testament (Psalms, Proverbs) and the New Testament (Gospels).

Some examples of C:

- (2) Lat. *corpus/corporis* 'body' [N, declension III] > Late Lat. *corpus/corpi* [M?, declension II] > OFr. *cors* [M] 'body'
- (3) Lat. *argentum* [N] 'silver' > OFr. *argent* [M] 'metal, money'
- (4) Lat. *vigilia* [Npl.] 'vigilance' > OFr. *veille* [F] 'wake, surveillance'

As regards subclasses of nouns, the change in gender is usually given by listing and is explained, if at all, by appealing to lexical semantics. For example, most names of trees and some other plants, which ended in *-us* but were feminine in Latin, became masculine, and the generic word *arbor* 'tree' is assumed to have then followed suit.

- (5) Tree names: *moorus* 'mulberry', *piinus* 'pine', *pirus* 'pear', *siicomorus* [*syy-comorus*] 'sycamore'; Lat. *arbor* [F] 'tree' > OFr. *arbre* [M]

Conversely, many abstract words with the accusative singular form in *-oor(em)* became feminine, and then the nonabstract 'flower' supposedly followed suit.

- (6) Feminines in *-oorem*: Lat. *dolor* [M] > OFr. *doleure* [F] 'pain, suffering'; Lat. *labor* [M] > OFr. *labor* [F] 'hard work'; Lat. *floos* [M] > OFr. *flor* [F] 'flower'

At the level of individual words, some nouns in Late Latin and in early Old French became or remained neuter before they stabilized as masculines or feminines in Old French of the later period. In early Old French texts, neuters constitute about 4.6% of the nominal vocabulary (masculines are at 48.3%, feminines at 47.1%). To illustrate, the Old Latin *gluuten* [M] 'glue' becomes *gluten* [N] in Late Latin and is attested as a neuter in Old French *gluz*; the Old Latin *gladius/gladium* 'sword', whose gender was either masculine or neuter, stabilizes as neuter in Late Latin (*gladium*), changing to *glai(e)* [F] in Old French.¹¹

The uncertainty of gender and/or vacillation of genders was part of the gender system in Latin and was carried over to Late Latin and early Old French. Some examples of the nouns of unstable gender in Late Latin include:

- (7) Unstable gender nouns in Late Latin: *cinus/cinis* [M/N/F] 'ashes', *dorsum/dorsus* [N/M] 'back', *cornum/cornus* [N/M] 'horn', *caput/capus* [N/M] 'head', *cyma (cuma)* [N/F] 'young cabbage sprouts', *latus* [N/M] 'side, flank'

Similar cases of gender vacillation are also attested in early Old French (see examples in Wartburg 1888, Godefroy 1901, Meyer-Lübke 1923, de Gorog 1982, Ayres-Bennett 1996:23):

- (8) Old French nouns of variable gender [M/F]: *art* 'art', *cantique* 'chant', *cervis* 'nape of the neck', *c(h)ifre* 'number', *dent* 'tooth', *font* 'fountain', *empire* 'control; empire', *isle* 'island', *levre* 'lip', *signe* 'sign'.

The word *mare* 'sea' changed from neuter to feminine putatively by analogy with *terra* 'land' (Meyer-Lübke 1923:56ff., Pope 1961:304–5). The change in the gender of 'milk' is not so easily explained. An early variant, Old Latin neuter *lacte* (vs. the more

¹¹ The situation with both of these words is not straightforward. For *gluten*, the *Oxford Latin dictionary* (*OLD*) lists it as neuter only, whereas Pope (1961:303) insists it was masculine in Old Latin; Ernout and Meillet (1939) also cite the form *glutis* which seems to appear later and fluctuates between masculine and feminine. Both *gladius* and *gladium* are attested in Classical Latin, but opinions differ as to which gender is older. For example, the *OLD* assumes that both neuter and masculine are equally old, while Lewis and Short (1962) take the neuter to be the older form. The situation is further complicated by the possibility that the word may be a Celtic loan into Latin (we would like to thank Brian Joseph for bringing these complicating factors to our attention).

usual *lac*), attested for example in Plautus, provides a basis for a reanalysis as if standing for *lactem*. Such a reanalysis may have been enhanced by the generally weak articulation of the final *-m*. Based on the putative form *lactem*, a masculine nominative *lactis* could have been formed; *lactis* is actually attested in Oribasius, a fourth century AD author (Ernout & Meillet 1939).

(9) Lat. *mare* [N] > OFr. *mare* [F] ‘sea’ ~ *terra* ‘land’ [F]

(10) Lat. *lac* [N] > Late Lat. *lact(em)* > OFr. *lait* [M]

2.4. MAIN REASONS FOR GENDER RESTRUCTURING. Over the history of Latin, the dissonance between form and gender gradually increased, and this was due primarily to changes in the shape of Latin words. A complete summary of the phonological changes that occurred over the history of Romance is certainly beyond the scope of this paper. For our purposes, the crucial changes included the loss of vowel length (beginning in Late Latin), rampant palatalization, and an almost complete loss of word-final consonants.

The loss of length distinctions leads to the correspondences shown in Table 5 between Latin and Proto-Western Romance, the precursor to French (Vincent 1988:32–33):

LATIN	ii	i	ee	e	aa	a	oo	o	uu	u
PWR	i	e	e	ɛ	a	a	o	ɔ	u	u

TABLE 5. Vowel correspondences between Latin and Proto-Western Romance.

Classical Latin already had a very limited set of word-final consonants: the nasals /m, n/, the liquids /r, l/, the fricative /s/, and the stops /k, t, d, b/; the last two were much less frequent (as evident from Diederich 1939) and hardly relevant to the nominal lexicon. Various changes affected this set, the most critical (and earliest) of which was the loss of word-final /m/, which brought chaos to the declension system. Recall that the final segments play a crucial role in determining the declension type of a noun, and this type in turn determines gender assignment. The loss of word-final distinctions inevitably led to a greater syncretism of case forms and genders.

An additional role in gender restructuring could be attributed to the Celtic substrate (Dottin 1920, Whatmough 1963, 1970, Schmidt 1967, Lejeune et al. 1985, Lejeune 1985, 1988). The Continental Celtic language spoken in Gaul had two genders, masculine and feminine. There were also some nouns that were neuter (in the *o*-stems). However, the scant evidence for the neuter class is not unequivocal and is hardly apparent from the nominal paradigm, as shown in Lambert 1994. In fact, whatever neuters existed in the Gaulish gender system of that time were very rare and formally hard to distinguish from the other genders, which further minimized their role in the system (Pedersen 1913:65–67, Lambert 1994:49, 50, 59, 60, 151, 157). Thus, for our purposes they can be ignored. The semantic gender assignment rules did not differ from those of Latin, but the rest of the nouns obeyed a completely different set of morphophonemic assignment rules. It is well established that the speech community of Gaul went through a period of protracted bilingualism (e.g. Wartburg 1888:65–67). This suggests that Celtic interference in Gallo-Roman could have been quite significant.¹² In particular, bilingualism could have affected gender assignment in both languages.¹³

¹² The formation of French was also affected by the German superstrate (Frankish); see Wartburg 1888 and McKitterick 1989, among many others. Although some scholars believe that Old French was strongly influenced by Frankish in a context of prolonged bilingualism (e.g. Malkiel 1975), this bilingualism was probably less widespread than the Gaulish-Latin bilingualism, and it was also more localized. Since we had to limit our simulations in order to proceed in a more practical manner, we chose to exclude the Germanic superstrate.

¹³ There are surprisingly few studies of the effects of coordinate bilingualism on gender assignment in an individual’s language. In her study of Italian-German bilingualism, Taeschner (1983:vii, 119–26) presents

The fact that the substrate had two genders, with just some traces of the neuter, may have facilitated the loss of the Latin neuter — the number of gender classes in the target language could have been assumed to be the same as in the substrate (see Carroll 1995: 201 for a general discussion of the decision that a speaker has to make with respect to the number of gender classes). As far as gender assignment of individual words by the bilingual speakers of Gaulish and Late Latin is concerned, there are a number of ways in which the two gender systems could have interacted, and little is known about the principles of such interaction (see n. 13). It would be too simplistic to predict that the only contribution of bilingualism is to achieve a match between genders in the two languages. In theory, the interference of one gender system in another can be based either on the similarity of concepts or on the superficial similarity of words. With each type of superficial similarity, the two languages spoken by a bilingual may either agree, enhancing the same motivation, or disagree, thus providing competing motivations in gender assignment. Thus, given concept *C* expressed by word *L* in Late Latin and word *G* in Gaulish, we can anticipate the following possibilities in the interaction between Latin and Gaulish (of course, noninteraction is still another possibility).

- A. the gender of *L* matches the gender of *G*;
- B. the gender of *L* is made different from the gender of *G* (on the implicit and subjective expectation in the bilingual mind that the two language systems should be dissimilar);
- C. the gender of *L* matches the gender of a phonetically similar noun in Gaulish (which expresses a concept other than *C*);
- D. the gender of *L* is made different from the gender of a phonetically similar noun in Gaulish.

In presenting these possibilities, we do not make any claims about which of them are more plausible than the others; to understand the interaction of all the possibilities, one has to either collect synchronic experimental evidence and/or conduct computer simulations. Our goal is to identify them all and to take into account as many as possible in the particular bilingual situation we are exploring, that is, Gaulish-Latin. The known Gaulish lexicon is quite restricted (Dottin 1920, Lejeune et al. 1985, Lambert 1994) and does not allow us to establish too many formal similarities that would be the basis of options C and D. The limitations of the data thus constrain the extent to which we can investigate the role of the substrate in the change from Late Latin to Old French. The general hypothesis we pursue here is that the choice of gender for a Late Latin noun could be influenced by the gender of the word expressing the same/similar concept in Gaulish, with the main evidence for this drawn only on the basis of possibilities A and B. As we show below, even this restricted evidence yields experimental results that are superior to the results based on a monolingual model, thus providing additional support for the use of computational models in reconstructing language history (see Kirby 1999 for similar observations based on a different set of simulations).

3. THE MODEL

3.1. CONNECTIONIST MODELING OF DIACHRONY. In the previous sections we have covered the historical data on the changes in the gender system between Late Latin and Old French. We now turn to the model which we have developed to account for these

some cases of gender interference in a child's speech. See also Grosjean 2000 on bilingualism in general and Guillemon & Grosjean 2000 on French-English bilinguals. The latter study, however, addresses a very different situation, namely, the interaction between a language with grammatical gender and a genderless language.

changes. Before we go into technical detail, though, it seems appropriate to motivate our use of neural network simulations to study historical linguistic data. Any such project is likely to raise two concerns. At a synchronic level, natural language is far too complex to be captured adequately in a relatively small computer model. From a diachronic perspective, artificially simulating a particular case of linguistic evolution—however successfully—does not necessarily provide any insight into the actual processes that drove the historical changes. Different processes can lead to results that are virtually indistinguishable, and there is often insufficient data available to decide conclusively between alternatives: for example, imagine how much easier it would be to account for the changes in the nominal system from Late Latin to Old French if we had access to longitudinal studies of contemporaneous children acquiring these languages and to speaker variation within the respective communities.

Because of the limitations, the goal of our model is not to provide the definitive account of the evolution of the gender system in Old French. Rather, the simulations are tools to investigate a more general issue—WHICH LINGUISTIC FACTORS ARE SUFFICIENT TO LEAD TO LANGUAGE CHANGE. We chose the nominal system in Late Latin and Old French largely because extensive philological data have been accumulated to document the respective periods in the history of Romance. While we personally believe that the historical changes from Latin to French were very likely guided by principles similar to the ones which guide our neural networks, the hypothesis we are presenting is more general and does not depend entirely on how well the simulations correspond to this instance of linguistic change. Our main hypothesis is that changes in noun class (gender) systems can be accounted for on the basis of a small number of simple semantic characteristics (e.g. natural gender) and of morphophonemic and frequency information which is often most ‘visible’ to a language speaker. The morphophonemic and frequency information plays the crucial role in our modeling of gender change.

Specifically, we predict that Latin neuters will be harder to learn for our models and will therefore be assimilated into the masculine or feminine classes on the basis of formal similarities. To restrict the formal cues available for the assignment of gender, we concentrate on word-final segments. This decision is motivated by the crucial role that word endings play in the identification of a declension class in Latin (see Tables 1 and 3).¹⁴ Given the protracted bilingualism of the speech community that implemented the transition from Late Latin to Old French, some changes in the gender system may be due to interference from the gender system of the substrate. Thus, we expect a network that has access to the substrate gender information to perform better than a network without such information.

3.2. CONNECTIONISM AND LANGUAGE HISTORY. There is a growing literature on the use of artificial neural network models for research into a wide range of linguistic phenomena. Much attention has been given to modeling patterns in language acquisition—the English past tense being the best known example (Rumelhart & McClelland 1986, Plunkett & Marchman 1991, 1993)—but there are also simulations that have examined data from morphology, syntax, language typology, and aphasia (Gasser 1993, Smith 1995, Elman 1990, 1992, Van Everbroeck 2003, Plaut 1999). Connectionist models also benefit from a larger degree of cognitive plausibility than many other

¹⁴ The role of IMPERIALISTIC ENDINGS has been addressed mainly in the acquisition literature under Slobin’s Operating Principle A: ‘pay attention to the ends of words’ (Slobin 1973, Smoczyńska 1985:667–68). Processing evidence suggests that hearers uniformly pay more attention to the initial part of a word (Allopenna et al. 1998). We expect both factors to play a role in gender assignment.

learning mechanisms. Recent work in neuroscience has shown that error-driven learning algorithms which are similar to the widely used backpropagation algorithm are indeed implemented by the brain (Fletcher et al. 2001).

Our connectionist model of the restructuring of Romance gender was inspired largely by the work by Hare and Elman (1995), who used sophisticated connectionist simulations to model changes in the past tense system of Old English. Among other historical developments, Hare and Elman were able to model the assimilation of weak verb class Ic (*nerjan*-type verbs) into weak verb class II (*lufian*-type verbs). This development bears a strong resemblance to the merger of gender classes in Old French, and a brief discussion of it will be illuminating here.

For a full discussion of the English past tense system and how it was modeled, the reader is referred to Hare & Elman 1995. Here, let us just mention that weak verb class Ic was distinguished from weak verb class II in two ways. First, verbs in class Ic took the derivative affix *-j-* between the stem and present tense suffixes, while verbs in class II took the affix *-i-*. Second, verbs in class Ic inserted *-e-* as the 'medial vowel' between the stem and the past tense suffixes, while verbs in class II inserted *-o-*. The first distinction was lost when glide vocalization changed the *-j-* affix of class Ic to *-i-*, the affix of class II. At this point, only the medial vowel distinguished the two classes. Hare and Elman hypothesize that 'as a result of this increased similarity, the two classes collapsed into one. The verbs of the small Ic class adopted the medial vowel *-o-* of class II, becoming indistinguishable from the original members of that class' (Hare & Elman 1995:68). They further speculate that the reason class Ic was absorbed by class II, rather than the reverse, was that class II had many members, with relatively open criteria for membership, whereas class Ic was quite small.

The similarity between these developments in Old English and the gender restructuring we are dealing with should immediately be clear. First, the neuter class in Late Latin was small compared to the masculine and the feminine (see Table 3). The second similarity between verb class changes in Old English and gender class changes in Late Latin is in the massive loss of phonological distinctions between separate classes. Latin gender assignment rules were morphological in nature. Gender assignment required access to more than one case form of a Latin noun — in other words, the assignment of gender was dependent on the declension class of the noun. As phonological distinctions between noun classes became obscure, some gender classes shifted into others. Unlike the English past tense, however, the Latin gender assignment rules were more pervasive (because the assignment of gender affected related clausal constituents, not just the lexical item itself), involved more distinctions, and had to respond to more subclasses of different size. In our choice of features for modeling, we tried to make the simulations linguistically sophisticated by representing the forms phonetically and by taking into account the morphophonemic similarities as well as the type and token frequencies of the nouns. Given the history of the Celtic substrate described above, we also introduced the gender of the corresponding Celtic noun. In doing so, we adopted the most general noun that expressed a given concept in Gaulish. To include the Celtic substrate in our simulations, we made a conscious decision to assign the gender of the interfering noun on the basis of the concept. Thus, we took into consideration the gender of the most general Celtic noun denoting the same concept as the Late Latin word. The general hypothesis is that the choice of gender for a Late Latin noun could be influenced by the gender of the word expressing the same/similar concept in Gaulish.

3.3. CORPORA. The nouns used to construct the training corpus for the network were selected on the basis of their frequency in the fifth-century Vulgate, excluding proper names and obvious Greek words of a liturgical/religious nature.¹⁵ The final list contained the 500 most frequent common nouns (see Appendix), as well as their raw frequencies. The latter, however, were not suitable for building the training corpus for two quite different reasons. The most important one was that equal raw frequency differences are not always equally important. Marcus and colleagues (1992) have argued that the frequency difference between 1,010 and 1,001 occurrences of a word is unlikely to be as relevant in any given corpus as the difference between 10 and 1 occurrences. In the first case, the difference is quite probably not statistically significant, and another corpus of the same size might show a similar difference in the opposite direction. But the difference between 10 and 1 occurrences is relatively speaking much greater and is also more likely to be relevant. For our simulations, therefore, we wanted to do justice to this latter kind of frequency difference, while simultaneously ignoring the less important differences. The second reason was much more practical in that our computers did not have the resources required to deal with the raw frequencies directly. The solution for both issues was to apply a log function to the raw frequency values and use those numbers instead (Marcus et al. 1992, Jackson & Cottrell 1997).

In the discussion of gender assignment above, we showed the correlation between gender and declension class. For most Latin nouns, the declension class can be derived from the form of the genitive singular (see Tables 1 and 4); for Late Latin, the relevant cases are nominative, genitive, and accusative, and these case forms were introduced into the network. Because we wanted the network to see an accurate sample of Late Latin, we next determined the relative frequency of each case and number combination (e.g. genitive singular, or accusative plural) in Late Latin texts. We also calculated separate statistics for [+Human] versus [-Human] nouns. The normalized results from a random sample of 1,000 sentences from the Vulgate and Plautus are shown in Table 6. For example, the nominative singular of [+Human] nouns occurred 8 times as frequently as the least frequent case/number combination (i.e. [-Human] genitive plural nouns); their accusative plural counterparts were only 2.5 times as frequent.

	[+Human]		[-Human]	
	SINGULAR	PLURAL	SINGULAR	PLURAL
Nominative	8	3	4	2
Accusative	4.5	2.5	7	2
Genitive	4	2	2	1

TABLE 6. Normalized distribution of [+Human] and [-Human] nouns across case forms used in the simulations.

By combining the log frequencies of the nouns and the values found in Table 6, it became possible to construct a training corpus which would accurately reflect how often a particular noun would appear with a particular case and number marker. To construct the total corpus, the token frequency of each of the six possible case and number combinations of the 500 nouns was calculated by multiplying the log frequency for each noun by the frequency of each of the combinations. The resulting value then determined the number of times the network would see each form of the noun in the training corpus. For example, *amicus* 'friend' had a log frequency of 8, so as a

¹⁵ These sources are obviously unlikely to reflect the precise word frequencies of the typical language input from caretakers and peers alike heard by children at the time, but we were limited to the texts that are still available.

[+Human] noun, its nominative singular form was put into the training corpus 64 times (8×8), whereas its genitive plural form *amiicorum* occurred only 16 times (8×2). The sum of all these noun forms ultimately was close to 38,000. The final step in the construction of the training corpus consisted of establishing the likely pronunciation of each of the noun forms.

The test corpus, which was used to gauge the performance of the network, was built using the same general procedure, except that each case and number combination of each noun was presented to the network only once. A single presentation of a form sufficed to determine whether the network classified it as masculine, feminine, or neuter, so there was no need to present a given form more than once. The total size of the test corpus was 2,918 forms. It was slightly less than the expected 3,000 (i.e. 500 nouns \times 6 forms) because a few nouns only occurred in either singular (e.g. *fides* 'trust') or plural (e.g. *intestiina* 'intestines') forms.

3.4. NETWORK ARCHITECTURE AND REPRESENTATIONS. For the purposes of our simulation, a simple feed-forward network with a single hidden layer was appropriate (Rumelhart & McClelland 1986).¹⁶ In this type of network, patterns are fed into the network at the input layer by setting the values of its units to either an active (1) or inactive (0) value—with an input layer three units wide; for example, 101 and 011 would represent different patterns. Active input units then excite or inhibit the units at the hidden layer with which they are connected. Excitation happens when the number associated with the connection between the units is positive, inhibition when the number is negative. Similarly, the units at the hidden layer send a wave of activation to the units at the output layer, and it is the pattern of activation at these last units which can be used to determine how the network classifies the original pattern fed into the input units. The architecture of the network used for our simulations is shown in Figure 1.

We may begin by taking a closer look at the two upper layers. The output layer of our network had only three units, one for each of the three possible genders. Only one of the three would be active for each input pattern during training, so that the network would learn to map this pattern onto, for example, the 'masculine' output unit. The hidden layer consisted of thirty units, all of which were connected to all the units in the input layer as well as the three output units. The purpose of this hidden layer was to give the network a chance to construct its own internal representations for finding similarities in the patterns presented to the input layer. So, this layer is where the network would construct abstract types and treat all input nouns that ended, for example, in *-us* as being quite similar irrespective of the phonological form of their roots. The hidden layer of the network is also where type frequency came into play, because the connections from the hidden layer to the output layer would be trained to map this *-us* type onto the masculine output unit much more often than onto the neuter output unit. So the neuter nouns in declension class IV (e.g. *corpus* 'body') were soon treated as if they were class II masculine words like *amiicus* 'friend'.

¹⁶ In recent years, recurrent networks like the ones pioneered by Elman (1990, 1992) have become quite popular for simulations of language. The main benefit of such networks is that the recurrent layer gives the model a short-term memory in which it can store information about the last few words that were presented at the input layer. While such an architecture is required for syntactic tasks like agreement detection, it would be needlessly complex for the current project: the task of our network was to classify each individual word, not to find patterns among sequences of words.

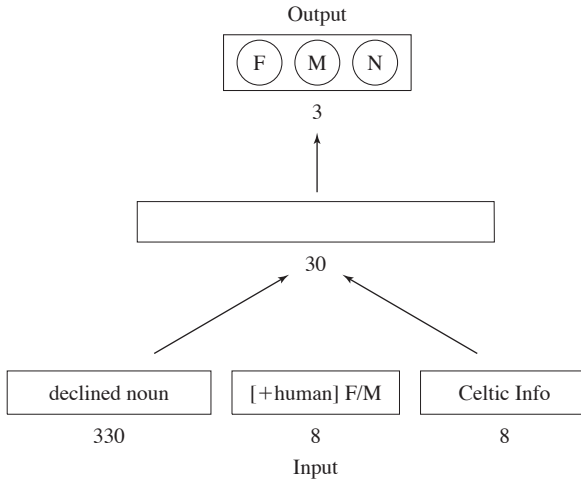


FIGURE 1. The architecture of the network used in the simulations.

The input layer consists of three groups of units to represent the phonological form of the word involved (330 units), its [+Human] and Late Latin gender status (8 units), and its gender in the Celtic substrate (8 units). The input layer feeds into a small hidden layer (30 units) where the input-to-output patterns are stored. The hidden layer in turn sends activation to the output layer (3 units) where the network can activate the output units to indicate the gender it assigns to the word shown at the input layer.

The input layer of the network consisted of three independent groups of units.

GROUP 1: PHONOLOGY. Each noun form could take up as many as five syllable slots, and each of these conformed to a template which accommodated at most six phonemes. Each of the phonemes was in turn presented as a sequence of eleven units representing phonological features in the general tradition of Chomsky & Halle 1968. If a particular feature was active for a phoneme, the value of that input unit was set to 0.9; if it was inactive, the input unit was given an activation value of -0.9 . Phoneme slots in a syllable which were not filled were given values of zero for all eleven feature units. Through multiplication, then, each noun was represented by 330 ($5 \times 6 \times 11$) units. The four examples below illustrate how the syllable and phoneme slots were filled for two forms of *aceetum* ‘vinegar’ and *adulter* ‘adulterer’.

```

----- --a--- k-ee-- t-u--m -----
----- --a--- k-ee-- t-ii-- -----
--a--- d-u--l t-e--- r-ii-- -----
--a--- d-u--l t-e--- r-oo-- r-u--m
    
```

Note that long vowels were distinguished from short vowels by doubling them. Note also the choice of syllable slots to be filled: the forms for *aceetum* occupy neither the initial nor the final syllable slot. Instead, its final syllable is in the penultimate syllable slot, as was the case for all noun forms except for those ending in a genitive plural *-oorum* or *-aarum*. The rationale for this representation is that experimental research has shown that people are aware of the identity of the endings of words (see above). Hence, native speakers of Latin would have noticed that *aceeti* (three syllables) and *adulteri* (four syllables) both ended in *-i*, and we wanted the network to have access to the same information. By filling the syllable slots in the manner illustrated above and reserving the final slot for the *-rum* part of *-oorum* and *-aarum*, we could ensure that the noun stems were aligned in the penultimate slot for all case and number combi-

nations and thus give the network a hint that this slot typically carried important gender information.¹⁷

GROUP 2: [\pm HUMAN]. These eight units were used to tell the network whether the word being shown in the phonology units referred to a [+Human] concept or not. If it did and the referent was [+Male], then the first four units were set to an active value 1 and the final four were set to 0. With a [+Female] referent these activation values were switched. If the referent was not human at all, all eight units remained at 0. The motivation for these units is that [+Human] referents usually appear to be more salient for human conceptual processes than their nonhuman counterparts. Also, people are aware of the commonalities in gender between words like 'girl' and 'maid'. The Late Latin counterparts to these words received added similarity through these units. Obviously, the correct amount of such similarity is hard to estimate, but the simultaneous activation of four units allowed the network to attach a lot of weight to this information if it wanted to. During training, it could increase the positive links between these units and those in the hidden layer, or decrease them into insignificance.

GROUP 3: CELTIC GENDER. These eight units were used in a similar fashion to the [\pm Human] units, except that the first four were made active if the Celtic counterpart to the Late Latin word was masculine, and the latter four if this counterpart was feminine. By giving the network four units each for the Celtic genders, we again wanted to give it the chance to make full use of the information if it considered it relevant.¹⁸

Since the notions of type and token frequency play an important role in the discussion section below, it is appropriate to describe their implementation in the model. The token frequency of a particular form in the model, such as genitive plural *amiicorum*, is the number of times it occurs in the training corpus (16). All the possible forms of a word add up to the token frequency for the word: for example, *amiicus* (nom. sg. 64) + *amiicum* (acc. sg. 36) + *amiicii* (gen. sg. 32) + *amiicii* (nom. pl. 24) + *amiicos* (acc. pl. 20) + *amiicorum* (gen. pl. 16) = 192. The type frequency of a pattern in the model is the sum of all the token frequencies that match it. Although there are many types of patterns (e.g. the type *amiic-* also occurs 192 times), the ones we are interested in are defined by the word ending. For example, the type frequency of *-us* as a pattern is 2,481; that of *-a* is 3,607. Finally, the type frequency of a mapping refers to the number of times a particular pattern is linked to a particular gender: for example, the type frequency of masculine *-us* (1,743) is much larger than that of either feminine *-us* (168) or neuter *-us* (570).

These frequencies are very important to the way the network works because they determine the number of times that a particular form, word, pattern, or mapping is trained on. With each presentation, the connections between the relevant input units and the hidden units are activated and then modified (during the error correction phase) depending on whether or not they led to the activation of the correct output unit, and similarly for the connections between the hidden units and the output units. So things that are more frequent are trained on more and are therefore learned better. The importance of type frequencies is also illustrated by cases in which many phonologically similar forms with low token frequencies together form a strong type — on occasion,

¹⁷ A similar representational problem was solved in a similar fashion by both MacWhinney and Leinbach (1991) and Ling (1994) in their models of the English past tense formation, despite the fact that the former used a connectionist network and the latter a rule-driven symbolic learning algorithm.

¹⁸ We did not consider gender information from other Celtic languages because we wanted to avoid possible discrepancies across the gender systems within the Celtic group.

such a type may even change the way a form with a much higher token frequency behaves.

3.5. TRAINING AND TESTING. Modeling the linguistic evolution of the Latin nominal system required that we adopt a connectionist mechanism capable of simulating language change. We chose to follow the same basic strategy first implemented by Hare and Elman (1995), in which generations of language learners are represented by generations of neural network models. Just as children begin their language development with the language spoken by their parents or other immediate caregivers, the model of a single network generation had to learn the ‘language’ used by that of the previous generation. And just as natural language speakers introduce some changes in their language, the models also display imperfect learning — in both cases, the next generation is therefore confronted with a different target language from the previous one. As time goes by, incremental changes generally create a new system which may be noticeably different from its linguistic ancestor. Something that is not captured by this model of language change is that adults also can innovate in the use of their own language. Although we do believe that children are the primary agents of linguistic change, our main reason for not including this factor in the simulations is that its implementation would have caused the desired target values for a network to change within a single generation; the effects of such a learning scenario in artificial neural networks are not well understood and would have made it much harder to interpret the models.

The procedure we used to create generations of networks is the same as in Hare & Elman 1995 and is illustrated in Figure 2.¹⁹ The first generation network is supposed to learn the perfect mapping from all its input words to their corresponding linguistic genders. However, the network falls short of that goal and ends up with a language that is close to the desired target but that nonetheless has some changes, that is, words that are assigned the wrong gender. These ‘mistakes’ happen because the error-correcting feedback that the network receives during training can be insufficient to overcome

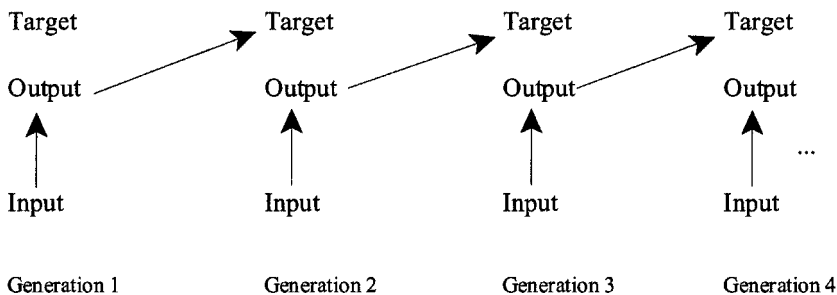


FIGURE 2. Procedure for modeling language change across generations.

We simulated generations of language learners by having neural networks try to learn the actual (imperfect) output of the preceding network. As each network generation falls somewhat short of the language of its ‘parent’, changes are introduced in the language. Over time, these small changes add up so that the language system of the last-generation network may look considerably different from that of the first generation.

¹⁹ Obviously, real societies cannot be divided easily into discrete generational groups, so this procedure to simulate generations is somewhat unrealistic. However, it was necessary to keep the simulations tractable. While it is possible to model populations of language learning networks where individuals would reproduce and die at varying rates (Batali 1998, Brighton 2002, Kirby 2002), no such simulations have ever been done with real, linguistically complex data such as ours.

the influence of other words with a similar phonological form but a different desired target gender. For example, the type that maps *-us* onto the male gender has such a high frequency that the network has a hard time learning that some words in *-us*, like *corpus* 'body', are actually neuter. The second-generation network is trained on the actual output of the first network, including its changes to the language. As one might expect, the second network will also display imperfect learning and introduce some modifications of its own. This process repeats itself over time until the networks have made enough changes to the original language that it can be acquired easily. We predicted that the neuter nouns would be the ones for which learning mistakes occurred most frequently, resulting in their being analyzed by the network as either masculine or feminine.

In our simulations, each generation was exactly three training epochs long — an epoch being a complete presentation of all the words in the training corpus — because we observed that this was sufficient for good overall training.²⁰ At this point, the exact response to each of the forms in the training corpus was saved. Then we used a simple Euclidean distance measure to 'clean up' the results: for example, if for any given form the trained network activated the feminine output unit strongly at 0.8, and the masculine and neuter units both at 0.1, then the training corpus for the next generation would have a desired output of 1.0 for feminine and 0 for the other two. The simplifying assumption here is that when speakers are presented with competing forms, they tend to pick the one that is most salient in their input and ignore the other forms as noise. Similarly, each new network started with a training corpus that had clear desired targets for each input form.²¹ At the end of each generation, the test corpus was also presented to the trained network to collect data on how the network was performing: that is, how many of the noun forms were given the original correct gender mapping, and how many were assigned to a different gender. These results are presented in the next section.

4. RESULTS AND DISCUSSION.

4.1. GENDER ASSIGNMENT IN GENERAL. To see how the generations of networks fared in general on the noun-to-gender mapping task, it is useful to look at the percentage of nouns, regardless of their gender, that the networks continued to map onto the original correct gender. This measure is not entirely unproblematic in that it continues to compare the output of every generation to the desired output for the initial generation, but it does allow us to assess fairly well how far the language produced by, for example, the fifth generation network has evolved from the original language. Recall that we are more interested in the total amount of change compared to the initial generation than in how much each new generation differs from the one just before it. We should also note that neural network models, like human beings, exhibit individual variation. In humans, such variation is the combined result of different genetic backgrounds and

²⁰ Given the nature of the problem, we did not want to train any network to perfect performance. It is very likely that a large enough network would manage to learn all the word-to-gender mappings given enough time. Children do not have access to infinite resources and infinite time, however, so the limitation on the number of epochs is not entirely artificial. All the simulations were implemented using the Stuttgart Neural Network Simulator 4.1 on Intel x86 Linux machines. We used a standard backpropagation algorithm for training with a small learning rate of 0.02. The code for generating and analyzing the various corpora and result files was written in a mix of C and Perl, and is available from the authors upon request.

²¹ Obviously, speakers of natural languages are sensitive to patterns of systematic variation in their linguistic environment and will typically learn to produce the right variants in the appropriate contexts. However, implementing this type of structured variation into our network model would only have complicated the simulations and slowed down the pattern of language change that we wanted to study.

environmental influences. In neural networks, on the other hand, such differences are caused by the initial values that are randomly assigned to each of the connections and units. Depending on the task with which a network is confronted, its initial configuration may make it easier for the network to learn the input-to-output mapping, or it may make it almost impossible. To ensure that an unusually favorable initial configuration did not arise by chance, we ran the simulation three times with different random initializations. The results for the three runs are shown in Figure 3. Variation among them was minor, and all three runs displayed the same overall evolution. Hence, all further discussion is based on the averages of three sets of simulations.




FIGURE 3. Overall performance on the three genders.

Percentage of nouns assigned the correct gender (i.e. the appropriate gender in Late Latin) over ten generations in three different networks. The three networks follow similar curves: they can learn about 80% of the genders in the initial 0th generation, but this number drops slowly to about 60% at the end of the 9th generation.

The 0th-generation network is the one that was trained on the corpus with the gender mappings from Late Latin. Figure 3 indicates that this network already had problems learning all the mappings: only 80% of the noun forms were assigned to the correct gender by the trained network. In the next four generations, this number drops to 63% but then remains almost constant until the end of the simulation (57%). Over the first five generations, the network therefore changed the language until it became more easily learnable.²² By the last generation, the network no longer has any problems learning the mappings found in the training corpus of the previous generation, so there is essentially no further linguistic evolution. This stabilization points out a difference between a network and a real language. In natural language, evolution is more constant, though certain subsystems will sometimes remain stable while others change. Presumably, sufficiently extensive changes in one subsystem will strain other subsystems,

²² Although the gender system as a whole becomes easier to learn from one generation to the next, the individual changes between different generations are not necessarily monotonic. Whether a particular gender mapping is difficult in any particular generation depends entirely on the nature of the other mappings that have to be learned at the same time. Consequently, a mapping that the network had a hard time learning in one generation may actually become easier in the next as a result of gender changes in other words. Similar phenomena have also been observed in natural languages: for example, the change from human to animate in some Bantu languages or the changes in the semantics of neuter in Konkani (Corbett 1991:248–59, 314).

forcing them, in turn, to evolve; with the influence of other languages and various cognitive/functional, pragmatic, and social factors taken into consideration, it is understandable why no real language has ever become stable. Our goal was to model a particular segment of language, and we do not equate our model with natural language. Thus, the stabilization achieved after nine generations should not be regarded as invalidating the simulation presented here.

4.2. BREAKDOWN BY GENDERS. The data in Fig. 3, however, do not shed any light on which of the genders the networks found more troublesome to learn. Looking at what happened to the nouns originally belonging to each of the three genders is more appropriate, and the next three figures do exactly that. Each of them shows the percentage of the noun forms that was assigned to which gender, so we can observe in which directions gender changes were most frequent. Figure 4 shows how these percentages evolved for the nouns that were originally masculine.




FIGURE 4. Evolution of masculine nouns.

Of the nouns that were masculine in Late Latin, the vast majority (roughly 85%) remained masculine over the course of nine generations. The remainder became feminine at some point, although a few were also temporarily neuter.

Although the number drops a bit during the first generations, about 85% of the nouns that were masculine in Latin remained so over the course of our simulation. About 5% of the noun forms switched to feminine in the initial network model, and this number then slowly increased to 15%. The number of masculine nouns that turned neuter, however, was much smaller, and they disappeared entirely by the last generation.

The main source of masculines changing into feminines is declension class III, where a variety of subtypes and the presence of all three genders could lead to a certain hesitation over gender. Recall that in Late Latin the accusative plural ends in *-s* in both class I and class III (Table 4); the high frequency of accusatives may have reinforced the similarity of the two forms. Declension class I was the class of feminine nouns, whereas the gender of class III nouns had to be established by appealing to additional information. Under these circumstances, the nouns whose gender assignment was unproblematic could serve as the attractor class for nouns whose gender was not apparent. Within class III, a particular group of nouns that changed from masculine to feminine were nouns in *-oorem* (acc. sg.), such as Lat. *color, coloorem* (masc.), OFr. *color/coleur*

(fem.) ‘color, temper’; and *furor, furoorem* (masc.), OFr. *furor/fureur* (fem.) ‘rage’; see also the examples in 6 above. In explaining this change, traditional grammars note that these nouns are abstract, thus forming a semantically homogenous group. This explanation is not without merit; crosslinguistically, abstract nouns are often (but not always) confined to a single gender, and in Latin abstract nouns indeed tended to be feminine. However, the corresponding nouns remained masculine in Spanish, Portuguese, and Italian. The geographical limitations of this change constitute a problem for the semantic explanation. In our experiment, the network had no access to the semantic feature ‘abstract’; nevertheless the relevant change still occurred. This suggests that the real attractor could be nonabstract feminine nouns in *-oorem* occurring at high frequency; two such nouns are Lat. *soror, sorooorem*, OFr. (*suer.*) *soro(u)r* ‘sister’ (training corpus frequency 138, log frequency in the Vulgate corpus 6) and Lat. *uxor, uxooorem*, OFr. *oissor/uxor* ‘spouse’ (training corpus frequency 138, log frequency in the Vulgate corpus 6); the latter, however, disappears quite early (Wagner 1974:211–12), so its role must have been less significant than that of *soror*. Note that this kind of explanation, based on straightforward formal analogy, is more plausible from the learnability standpoint (a young child accurately acquiring genders is not likely to have a full grasp of abstract concepts) and is also more parsimonious. It appeals only to analogy, whereas the more traditional approach taken by historical grammars appeals to both analogy and abstractness in order to account for the feminine gender of OFr. *flor* ‘flower’ (from Lat. *floos, flooris*, masc.).

Some masculine nouns briefly changed to neuter. This may be a surprising result, given the expectation that neuter was to be absorbed by masculine and feminine nouns. However, this result meshes well with the historical data. In Late Latin and subsequently in Gallo-Romance, a number of nouns fluctuated between neuter and masculine, and some nouns did indeed change from masculine to neuter, as illustrated in 7 above.

Turning to Figure 5, we find that the results for nouns that were originally feminine are quite similar to those for the masculine nouns. Although the percentage decreases slowly over generations, the percentage of feminine nouns that remain feminine is always at least 85%. And as before, there are a number of nouns that are mapped onto the masculine output unit (from 4% to 15%), whereas there are only a handful of nouns that briefly switch to neuter (0.5% in the 0th generation).

These results again mesh well with the historical data. The most conspicuous class of nouns that changed from feminine to masculine were names of trees ending in *-us*, such as *piinus* ‘pine’, *cedrus* ‘cedar’, and *siicomorus* ‘sycamore’. This can easily be explained by assimilation to declension class II, which had many nouns in *-us* and served as a strong attractor. The generic noun *arbor, arborem* ‘tree’ also changed to masculine (OFr. *arbor*). Traditional grammars account for this change by appealing to the semantic association with all the other tree names. However, the change in the network occurs without such a semantic association. It is more plausible that *arbor* changed along with other nouns in *-or* that became masculine. Almost all such nouns started out as neuters, for example, Lat. *cor* > OFr. *coeur* ‘heart’, Lat. *marmor* > OFr. *marmar* ‘marble’. It is important that quite a few [+Human] nouns ending in *-or* were masculine. Such nouns, due to their type and token frequency, could have formed the formal attractor group for other nouns in *-or*, regardless of their semantic features. Again, we can account for the attested historical change by appealing to formal analogy and to the frequency of the potential attractor nouns.

Finally, we turn to the nouns that started out as neuters in Late Latin. Given that about 85% of the masculine and feminine nouns retained their original gender, we can

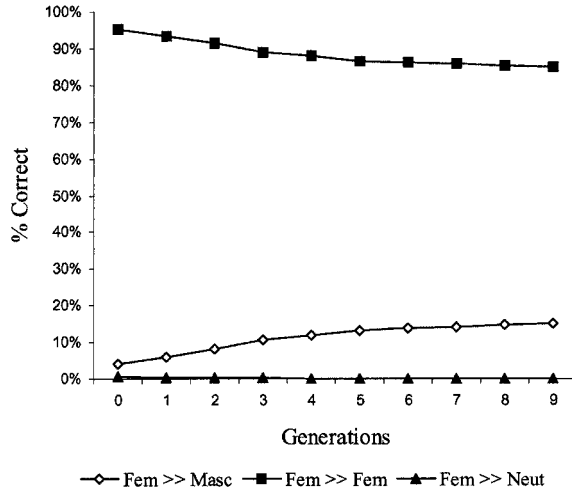


FIGURE 5. Evolution of feminine nouns.

Of the nouns that were feminine in Late Latin, the vast majority (more than 85%) retained their gender over the course of the simulations. The others became masculine, although a few were briefly neuter.

expect that the neuters were the nouns most affected in the simulation. This expectation is borne out by the evolution shown in Figure 6. The network already found the neuter nouns much harder to learn in the 0th generation, with only 51% of them mapped onto the neuter gender at the end of training. During the next nine generations, the percentage continues to decrease until it reaches 1% in the last generation.




FIGURE 6. Evolution of neuter nouns.

The nouns that were neuter in Late Latin slowly became either masculine (65% at the end of the simulations) or feminine (35%). Although the feminine gender was initially as strong an attractor as the masculine, the latter became dominant after the first generation.

Initially, the feminine nouns exert a slightly stronger pull as neuter plurals in *-a* are drawn into the feminine *-a* class: the network learns early on that a final *-a* is a good indicator of the feminine gender. Interestingly, the other case and number combinations

of the same nouns typically resisted switching much longer. So after only a single generation, the number of neuter plurals in *-a* that had become feminine was almost ten times larger than the corresponding number of singular neuters. In actual linguistic history, the development in which neuters with a plural in *-a* became feminine indeed began in Late Latin, so the early peaking of the neuter-to-feminine switch in the simulation is a satisfying result. The masculine influence catches up with the feminine nouns over time.

The main class of masculine nouns that incorporated neuters was declension class II. The historical assimilation of neuters to the masculines was due to the loss of the final consonant, which made the masculines in *-us* and the neuters in *-um* indistinguishable. But as the results of the model illustrate, this change might have occurred even in the absence of sound changes, being driven instead by frequency effects. Two other classes of neuter nouns that became masculine in the model were those ending in *-or* and in *-men*. Examples of the former are *ador* ‘spelt’ and *marmor* ‘marble’. The latter is illustrated by Lat. *crimen* > OFr. *crime* ‘crime’ and Lat. *noomen* > OFr. *nom* ‘name’, both of which switched their neuter gender to masculine both in history and in the model. Still, why did the neuters assimilate to the masculines and not the other way around? We believe the answer is type frequency and, related to that, learnability: masculine nouns were more numerous and thus easier to learn, which enabled them to absorb the recessive neuters.²³

Some nouns left the neuter gender class but vacillated a while between feminine and masculine; one such noun is the OFr. *joī(e)* ‘joy’ [M/F] (cf. Modern French *joie*, fem.), from the Latin neuter *gaudium* ‘joy’ (pl. *gaudia*). The plural form was assimilated to the feminine singular, but the singular was assimilated to the masculine in *-us*, and the tension between the two forms is responsible for the hesitation over the gender. Note that the relative frequency of the singular vs. the plural for an abstract noun like ‘joy’ will probably favor the singular form, which can account for the gender doublet.

4.3. EVIDENCE FOR THE CELTIC SUBSTRATE. In our simulations, we introduced into the network the data on Gaulish gender on the assumption that the Celtic substrate interfered with gender assignment in the Latin spoken by inhabitants of Gaul. If our assumption is correct, we should expect that similar simulations without the information on Celtic gender should produce results that do not conform as closely to the historical reality. The expectation is borne out by the data shown in Figure 7 (compare with Figs. 4, 5, and 6).²⁴ The difference for the masculine and feminine nouns is small when compared to the earlier simulations, but it is obvious that the neuter gender persists much longer than it does in the models that had access to the Celtic gender information. Even after training for more generations, the neuters never completely disappear.

Thus, the change in gender reassignment without the Celtic substrate shows less conformity to the facts of language history. This in turn confirms the hypothesis that interference from Celtic gender played a role in gender restructuring. Aside from histori-

²³ In a development peculiar to Gallo-Romance, the masculines of declension class III assimilated in the plural to the masculine plural of declension class II, for example, Lat. *pons, pontis* ‘bridge’ > **ponti* > OFr. *pont, ponts*. We could not simulate such a development in our model because the output layer of the network did not code which form the net thought each word should have (but see below).

²⁴ The data in Fig. 7 are again based on the averages of three simulations using different random initializations. The only difference from the simulations presented earlier is that the eight input units that could signal the gender of a word in the Celtic substrate were now always set to 0, irrespective of the Celtic gender.

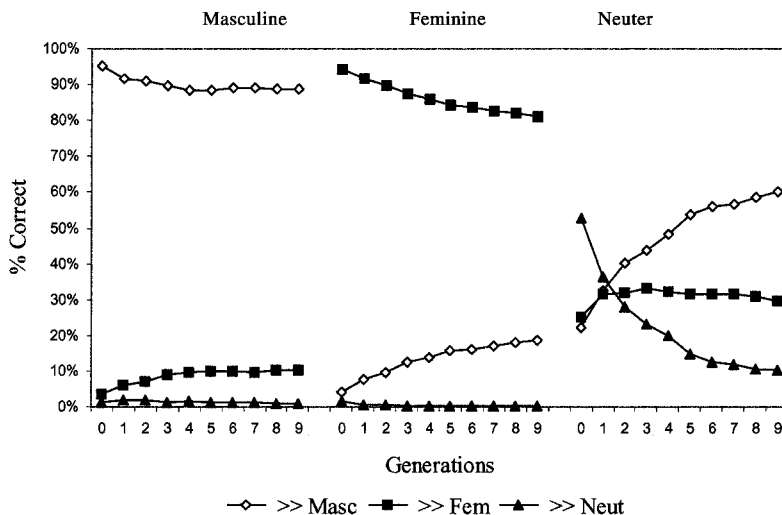


FIGURE 7. Evolution of nouns (without Celtic substrate).

When the gender information carried by the Celtic substrate was excluded, the nouns that were neuter in Late Latin retained their gender much better than when that information was included. About 10% of these nouns remained neuter at the end of the last generation.

cal implications, this result points to a more general conclusion, namely, that bilingualism produces gender interference.

4.4. RATE OF CHANGE IN RELATION TO WORD FREQUENCY. Our simulation of gender change relied crucially on the frequencies of individual words and frequencies of word classes (declension classes, genders, case forms). In the simulations above, we allowed all frequencies, drawing no distinction between high-frequency nouns, mid-frequency nouns, and nouns of lower frequency. Our corpus did not include very many words with low absolute frequency, so it makes sense to establish only relative distinctions within our list. We would now like to speculate on the differences between high-frequency words and all other words. High-frequency words have been shown to resist change more than low-frequency items (e.g. Bybee 1985:47ff.). Our results do not contradict this finding, but they do not add anything to it either. We do observe that once high-frequency nouns are excluded, the network actually converts neuter nouns into feminine and especially masculine ones somewhat faster than in the simulation including the high-frequency nouns. If we compare Figure 8 with Fig. 6, we find that removing the high-frequency nouns yields overall results that are more in conformity with the actual gender change from Late Latin to Old French.²⁵

If high-frequency words resist change, does this mean that low-frequency words should be more susceptible to change? Excluding high-frequency nouns from consideration, we found no difference between the rate at which low-frequency words and mid-frequency words reassign gender. Thus, we cannot confirm the strong prediction that low-frequency words should be most susceptible to change, just the weaker prediction

²⁵ The data in this figure were calculated by excluding all the nouns with a log frequency of 8 or higher (range 1 to 11) from the training and test corpora. This cut-off point entailed removing 58 of the 500 nouns. Of these nouns, 23 were feminine, 21 masculine, and only 14 neuter. Three simulations were again done with the remaining nouns, and the averages of their results are shown in Fig. 8.

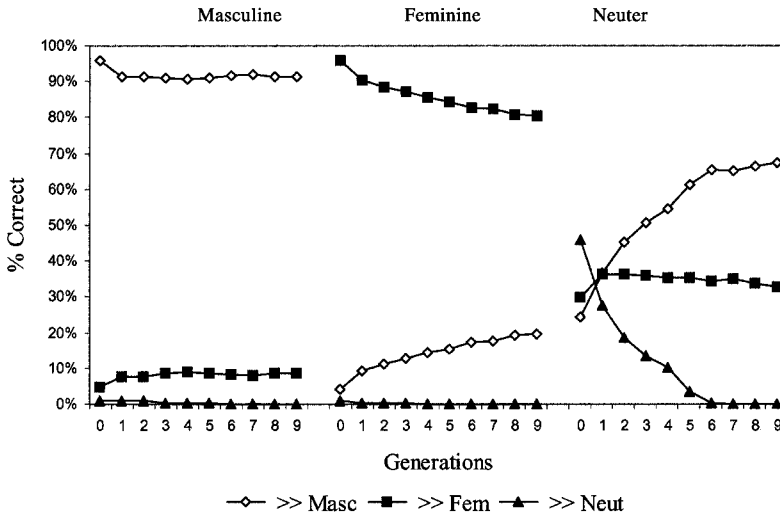


FIGURE 8. Evolution of nouns (without most frequent ones).

When the 10% most frequent nouns are excluded from the training corpus, the network reassigns the gender of the Late Latin neuter nouns to either masculine or feminine somewhat faster than in the original condition.

that high-frequency words are often stable. However, we do show that smaller classes of nouns are subject to faster change than larger classes — this is obvious from the early absorption of declension classes IV and V into the larger classes II and I, and from the partial assimilation of class III masculines to class II.

Another assumption that we used in interpreting our results has to do with the role of attractor words or attractor groups: namely, we assumed that high-frequency words provided a formal analogical model for words whose gender was uncertain and caused hesitation. This assumption, however, contradicts the hypothesis that high-frequency words have less effect on members of their classes than medium-frequency words do. According to this hypothesis, high-frequency words are more autonomous in every way, including being more autonomous from the classes to which they appear to belong (Moder 1992, Bybee 1995).²⁶ In our simulations, we observed that higher frequency items served as possible attractors for words with lower frequency. However, there may be at least two reasons for this difference between our results and the findings reported by Bybee and Moder. First, they investigated immediate responses to probes in a restricted experimental setting, where the subjects' attention was focused on isolated lexical items and they had to choose among nonce lexical items. In the restructuring of gender, the interaction of factors is more complicated, there is no need to choose one word over another, and the data are noisier, which makes comparing the two settings questionable. Second, the attractor words in our simulations were not necessarily highest in frequency across the entire corpus of words, which again distinguishes them from the absolutely high-frequency items used by Bybee and Moder. A better understanding of fine-grained distinctions within the overall category of frequency may be needed to

²⁶ This hypothesis was tested experimentally. In a nonce-probe task, Moder (1992) primed subjects with real verbs belonging to English strong verb classes and then presented them with nonce verbs similar in shape to the prototype for these classes. High-frequency primes had no effect, but mid-frequency primes had the strongest effect.

resolve the difference in our results; in particular, the distinction between absolute and relative frequency may become crucial in that relatively infrequent nouns could still serve as strong attractors for a small subgroup of words. In our simulations such a case is represented by the word *soror* which, as we hypothesized, served as a feminine attractor for a number of nouns ending in *-or*.

4.5. ANALYSIS OF ERRORS. The network failed to change the gender of a number of words, contrary to the historical data. In this section, we concentrate on classes of errors, rather than individual words.

A subclass of nouns that did not change in a uniform direction and that showed fluctuation is nouns in *-ns* (e.g. *dens* ‘tooth’, *fons* ‘fount’, *serpens* ‘serpent’, *oriens* ‘east’, *mons* ‘mountain’, *providens* ‘provident’). Some of these nouns were reanalyzed by the network as feminine (*dens*, *fons*, *mons*, *oriens*), the others as masculine, such as *fons* ‘front’.²⁷ However, this is a ‘welcome’ case where the inaccuracy in the results reflects an actual historical fact; quite a few nouns of this class did show fluctuation in gender between the masculine and the feminine in Old French (see 8), and there seems to have been a fair amount of uncertainty with respect to these nouns even in Vulgar Latin (Wartburg 1888:65, Elcock 1960:59–60, Pope 1961:304–5).

Another set of errors that can be explained post-hoc are errors in the gender assignment of imparisyllabic nouns — nouns with more syllables in the oblique cases than in the nominative (e.g. *boos*, *bovis* ‘ox’; *lux*, *luucis* ‘light’; *nix*, *nivis* ‘snow’; *mors*, *mortis* ‘death’; *urbs*, *urbis* ‘city’). We introduced these nouns into the network without any correction for the discrepancy in the number of syllables; thus, the network was faced with true imparisyllabic paradigms. In the meantime, it is a well-known historical fact that the imparisyllabic type was often simplified in such a way that the nominative (the shorter form) was reanalyzed to conform to the oblique stem; thus, *urbs* becomes *urbis* in the Merovingian period, and *pees* ‘foot’ is reanalyzed as *pedes/pedis* in Gallo-Romance (Pope 1961:309). Thus the failure of the network to reassign gender of imparisyllabic nouns may accurately reflect the fact that paradigm leveling (which we did not introduce into the model) preceded gender reanalysis. Among the imparisyllabic nouns, the nouns that uniformly and incorrectly changed from the feminine to the masculine in the network simulation are abstract nouns ending in *-io(n)* (*cognaatio(onis)* ‘relationship’, *possessio(onis)* ‘possession’) and abstract nouns in *uudoo/ -uudinis* (*pleenituum* ‘completeness’).

A number of errors could have been avoided if the simulation had included step-by-step phonological changes. For example, the nouns of declension class V remained feminine, but largely because they had been formally assimilated to class I singulars (*facies* > *facia* ‘form, shape’; *species* > *specia* ‘sight’; *dies* > *dia* ‘day’, hence OFr. *dia domenica* > *dimanche* ‘Sunday’). Without such assimilation the network cannot be expected to produce a historically adequate gender reanalysis.

Finally, some nouns have to be admitted as plain errors. For example, *lacus* ‘lake’ and *paanis* ‘bread’ were incorrectly reanalyzed by the network from masculine to feminine; *humus* ‘soil’ and *pellis* ‘leather, skin’ were incorrectly reanalyzed from feminine to masculine. We do not have a plausible explanation for these errors and have to admit it as a problem in the simulation. We would like to add, however, that the assumption that all change is principled is overly restrictive; random or stochastic events with a variety of causal factors ultimately leading to linguistic changes certainly occur

²⁷ This particular change accords with the historical data (OFr. *fronz/front* ‘front’).

in the history of languages and play a role in synchronic variation. A possibility that simulations reflect a certain degree of arbitrariness in language cannot therefore be excluded.

In addition to the historically incorrect gender assignment of specific nouns, a more general methodological problem is the difficulty of interpreting the results of the model when it gives conflicting gender changes for different case/number combinations of a single word. One solution might be to force a consistent gender for all case and number combinations of every word at the end of each generation; we could use a majority vote scheme with a random choice if the choices are equally frequent. However, this would be a decision that would conflict with the actual historical change—recall that in the reanalysis of the neuters in Late Latin the majority of singular neuters changed to the masculine and the majority of the plurals changed to the feminine (see §2.2 above); the result was a large class of nouns that were derivationally, but no longer inflectionally, related. If a forced choice were imposed in the interpretation of modeling results, this effect would be impossible to capture.

The general conclusion one can reach here is that the network adequately captured the initial stage of the three-gender system and the historical developments that were to follow. The initial system proved learnable, but the neuter class was already less accessible to error-free learning. We attribute this neuter handicap to the relatively small size of the class. Our network reassigned the neuters to the feminine and masculine classes, in conformity with the historical developments. Thus, plural neuters were reassigned to the feminine class, and this reassignment preceded the reassignment of neuters to the masculine class, because the latter was contingent on incremental phonological change. The reassignment in the network could be based only on formal analogy and frequency data, since no other information was available. In addition to the global restructuring of neuters, the network also reassigned some feminine nouns to the masculine class (see Fig. 5) and underwent a slight increase in neuters, also consistent with historical developments. The ability of connectionist simulations to represent historical developments has been known since the work of Hare and Elman. However, the simulation described here is the first to show the actual role of formal analogy as well as type and token frequency in the reanalysis of Romance genders. It also supports our claim that word endings played a crucial role in gender assignment.

5. MORE COMPLEX SIMULATIONS. One criticism that might be raised against the model presented here is that the kind of task that the network has to learn—mappings from noun forms onto one of three grammatical gender units—is not likely to be a close approximation of what native speakers are engaged in when they use nouns. In order to address this issue, we are currently working on a more complex model in which the networks have to map noun forms onto other noun forms. Instead of an output layer with only three units, the new model represents the six phonemes potentially occurring in any possible nominal case suffix. The suffix thus consists of sixty-six units—eleven for each of six phoneme positions in the suffix. At the input layer, a new group of units has also been added so that the network can be instructed as to which case/number suffix of the noun has to be produced at the output layer. So, instead of mapping a noun like *amicus* ‘friend’ onto the masculine output unit, the new model might have the phonological representation of *amicus* activated as well as units standing for ‘accusative plural’. At the output layer, the new model would have to produce the phonological representation of the appropriate suffix, that is, *-oos*. This architecture allows us to train the network to map all possible case and number combinations for a noun onto

the same set: for example, the training corpus might also contain the mapping from *amiicoos* and 'nominative singular' onto *-us*. This task seems more plausible than the original one in that speakers of a language also have to vary the case and number suffixes of nouns they use in discourse. The architecture of the new model is shown in Figure 9.

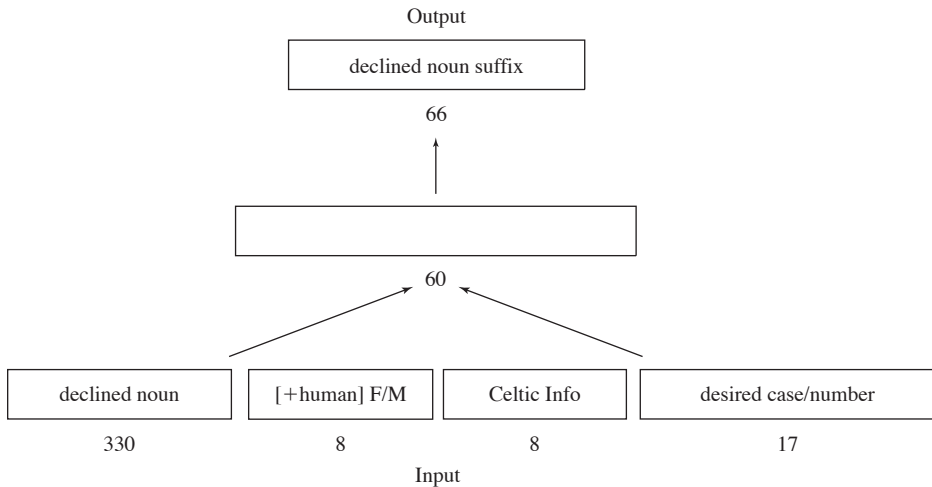


FIGURE 9. The architecture of the new model.

The most important change is at the output layer, where the network now produces the phonological representation of the appropriate noun suffix (as indicated by the 'desired case/number' group of units at the input layer).

The results obtained so far with the new model are promising, but the larger and more complex network, as well as the increased complexity of the task, makes it much harder to experiment with. One interesting phenomenon that we have already observed in some of our preliminary simulations with the more complex model concerns a compensatory effect. A network similar to the one in Fig. 9 was able to learn the desired mappings quite well; however, when we tried to make it learn the same mappings after applying the historical sound changes to the corpora, we found that it was no longer able to do so. This agrees with the historical hypothesis that the sound changes which took place between Late Latin and Old French made the gender system more opaque, thereby forcing speakers to learn more forms by rote. The resulting opacity led to more errors and more changes. One of these changes was the disappearance of the genitive case, resulting in a system with just a nominative and an oblique (formerly accusative) case. Somewhat to our surprise, we observed that if we removed the genitive forms from the simulations with the post-sound-change corpus, the network was again able to learn the mappings to almost the same level of performance as the three-gender system before the sound changes. We still need to confirm that this is indeed a robust effect, but it definitely suggests that neural networks could be used to model language changes in which increasing complexity in one part of the language forces compensatory adaptations in one or more other parts.

6. CONCLUSION. We have applied a frequency-based neural network, designed to approximate generational learning, to the learning of gender classification in Latin. The network had access to basic semantic information relevant for gender assignment in

Latin (natural gender of human referents) and to the declension class of each noun, which, to a large extent, determines the gender of that noun. Declension class is determined on the basis of several case forms, which were presented to the network with their actual frequencies. Thus, aside from the superordinate semantic information, the network had access to formal cues (case endings) and type and token frequencies of case forms, declension classes, and subsequently genders. This information proved sufficient for the adequate learning of genders in Latin; errors typically occurred in those paradigms where gender assignment was ambiguous and where speakers of Late Latin actually showed hesitation and made errors. The initial simulations therefore confirmed the presence of a heavily analogical system based on formal cues.

It has been shown (Plunkett & Marchman 1991, 1993) that multiple generalizations can create tensions in analogically-based systems, and the internal tension is later relieved by the restructuring of the system. In the transition from Latin to Old French, the major push for the restructuring of the gender system came from phonological changes (loss of vowel length, loss of word-final segments). The phonological changes led to numerous ambiguities, which made the earlier system much harder to learn and which led the smaller classes to assimilate to larger classes on the basis of formal similarity. In particular, a number of plural neuters assimilated to singular feminines based on the identity of endings (in Latin, both groups shared the ending in *-a*; in Old French, the common ending was *-e*). Meanwhile, a number of singular neuters assimilated to the singular masculine. In addition to this major restructuring, the network adequately captured smaller changes, such as a subset of feminine nouns changing to masculine, a subset of masculine nouns changing to feminine, and the rise of 'spurious' neuters. In smaller subsets of words, we hypothesize that a single word or a small group of words of higher frequency could have acted as an attractor for the lower frequency words. This conclusion, which is based on the results of our simulations, leads to the interesting theoretical question of whether or not a small group of words can play a role in triggering analogical change. We feel that this result needs to be tested in experimental work that would allow one to control for various factors that may only be hypothesized about in historical studies.

Many speakers of Late Latin were bilingual in Latin and Gaulish (Celtic). We hypothesized that the Celtic substrate, in which only a tiny proportion of nouns belonged to the neuter class, could have been an additional factor in the speakers' partitioning of the nominal lexicon into two, not three, classes. Additionally, the gender assignment in the substrate could have interfered with the gender assignment in Latin. The simulations suggest that this hypothesis is on the right track — the output of generational learning without the gender information from the substrate is much less historically accurate. We did not present the actual Celtic words to the network, which suggests that the abstract gender, even when associated with a given concept and not with a concrete lexical item, can play a role in the bilingual system. This calls for future experimental work on the interaction of gender assignment in bilinguals; the simulations per se cannot shed light on the facilitation vs. inhibition of gender assignment based on the substrate information.

Latin-to-Old French gender reanalysis has traditionally been accounted for by appealing to the interplay of formal and semantic factors. Although we limited the semantic information to natural gender of human referents, our simulations achieve high descriptive and explanatory adequacy. With respect to descriptive adequacy, the network follows most of the attested historical developments. As for explanatory adequacy, the results confirm that the Latin-to-Old French gender reassignment was based on a small

set of phonological and morphological cues (natural gender in humans, of course, also plays a role). The results also provide us with an important insight into the role played by frequency in formal analogy, suggesting that frequency effects are more complex than is often assumed. In particular, high frequency items are ensconced in small classes of their own and do not influence the overall analogical patterns observed in diachronic processes.

Our model is a further step in the development of complex computational networks of actual language change. Its main advantage over the model presented in Hare & Elman 1995 is that our model is closer to using true phonological representations at the input layer, which makes it easier to establish and manipulate the forms that the network is considering. Of course, we are fully aware of the limited nature of the tasks faced by the networks presented here, and we certainly do not pretend to have solved all the problems of gender assignment.

In sum, we have been able to account for change in gender between Late Latin and Old French using a sophisticated connectionist model that had access to superordinate semantics and formal cues. The model also showed that taking Celtic gender assignment into account significantly improved the fit with the historical data. Overall, these results show the utility of connectionist modeling in explaining language change. We hope that this discovery tool will be recognized and used in a broader context. At least three possible applications come to mind: the testing of historical reconstructions where only the final stage is known, the modeling of gender processing in bilinguals, and the ranking of synchronic versus diachronic factors in gender assignment.

APPENDIX

The list of 500 nouns used in the simulations is given below. Nouns marked as (sg) occur only in singular forms; those marked as (pl) occur only in the plural.

Masculine nouns: aaeer, adulter, aestus, ager, agnus, agricola, amiicus, amor, angelus, angulus, annus, ariees, baculus, boos, caasus, cadaaver, calix, cameelus, campus, canis, capillus, carboo, carcer, census, centurioo, chorus, cinis, circuitus, circus, claamor, clangor, coetus, collis, color, conspectus, convictus, craateer, cucumis, cultus, cuneus, currus, cursus, decor, deebitor, dens, deus, diabolus, di(i)es, discipulus, dolor, dominus, dracoo, dux, equus, exercitus, fiilius, fiinis, floos, fluvius, fluxus, follis, fons, fraater, fremor, fundus, furnus, fuuror, fuuniculus, fuunis, gladius, gradus, grex, gustuus, gyyrus, homoo, hospes, hostis, hyacinthus, hypocrita, ignis, inimicus, laabor, lacus, languor, lapis, la(a)troo, lector, lectus, libellus, liber, liiber, liimes, liivor, locus, lumbus, lupus, magister, magus, modus, mons, mundus, muurus, naasus, nepoos, numerus, obitus, obolus, occidens, oculus, odor, ooraator, orbis, ordo, oriens, paanis, paar, parens, passus, pastor, pater, pauper, pees, pepoo(n), percussor, peregrinns, persoona, pilus, piscaator, piscis, pontus, populus, porcus, porticus, praeceptor, princeps, proocessus, proovidens, propheeta, psalmus, puer, pullus, pulvis, pusillus, raamex, raamus, reex, riitus, rusticus, sacerdoos, sanctor, sanguis, scopulus, seemis, senior, sensus, sermoo, serpens, servus, sexus, siclus, somnus, sool, specus, spiiiritus, suucus, terminus, thronus, tinctus, torrens, tribus, tribuunus, turtur, tyrannus, unctus, unguis, uterus, vectis, vectus, venter, ventus, vermis, versus, vertex, vesper, victor, viiciinus, vir, vultus.

Feminine nouns: aarea, abyssus, adulter, adustioo, advena, aestaas, aetaas, ancilla, anima, aqua, aquila, arbor, arca, areena, asina, barba, basilica, basis, bestia, caaliigoo, capra, causa, cedrus, cicoonia, ciivitas, cognaatioo, columba, columna, commootioo, coniuga, coroona, cortiina, crusta, crux, culpa, custoodia, cuura, doctriina, domina, domus, ecclesia, epiistola, exaltaatioo, faama, faabula, facies (sg), fallaacia, fames (sg), familia, famula, fariina, feemina, festuua, fides (sg), figuura, fiicus, fiilia, fiscella, flamma, forma, fores (pl), frons, frux, fuga, gens, glooria, graatia, gula, herba, hoora, hospes, hostis, humus, hypocrita, iira, industria, injuuria, iustitia, laana, la(a)crima, lacuuna, lampada, lex, ligaatuura, lingua, lucerna, luucusta, luuna, lux, maater, maateria, magister, malitia, manus, margariita, memoria, mensa, mensuura, merces, meridiies (sg), miica, mola, mora, mors, mulier, multituudoo, murra, musca, naavis, negotioo, nivees (pl), nox, nurus, nuubees, oora, opera, opiinio, ovis, paala, palea, palma, parabola, parens, pars, pa(a)tria, pauper, pax, pecus, pecuunia, pellicula, pellis, pelvis, pestis, pe(e)tra, piinea, piinus, plaga, platea, plebs, pleenituudoo, pluvia, poena, porta, possessioo, potestaas, praeda, primitia, puella, pugna, raadix, ratioo, rees, regioo, religioo, rosa, ruiina, sapientia, scientia, scoopa, seemita, sententia, series (sg), societaas, socrus, soolituudoo, soror, species

(sg), spiina, sporta, stella, substantia, summa, suubula, syycomorus, temptaatioo, tene(e)brae (pl), terra, testa, torta, trabs, tuba, turba, turma, turris, tuuteela, umbra, ungula, urbs, uuva, uxor, vacca, vallis, verna, vestis, via, victooria, viiciniia, viinea, viipera, viis, viisioo, viita, villa, virga, virgoo, voluntaas, vox, vulpes, vulva, zoonaa.

Neuter nouns: aatrium, absinthium, aceetum, ador, altaaria (pl), animal, argentum, arma (pl), aurum, baasium, baptismum, bellum, bonum, bra(a)chium, caelum, canistrum, caput, carmen, castella (pl), castra (pl), concilium, cor, cornuu, corpus, crimen, cubiculum, daemonia (pl), deeserta (pl), exemplum, factum, ferrum, festum, fluumen, foraaamen, forum, fragmentum, fretum, furtum, gaudium (pl), genus, graanaata (pl), horrea (pl), incensum, infernum, intestiina (pl), jugum, juumentum, juus, labium, lac, laganum, lignum, luumen, malum, mare, marmor, medium, mel, membrum, milia (pl), miiraaculum, monumentum, muunus, noomen, officia (pl), oleum, oliiveetum, oppidum, os, ostium, pallium, peccaatum, pectora (pl), pectus, pinnaaculum, plaustra (pl), poomum (pl), pondus, pretium, reete, saeculum, scandalum, seecreetum (pl), seemen, signum, speculum, talenta (pl), tectum, templum, tempus, tergum, uuniversum, vaas, veelum, verbum, verum, vestiimentum, vexillum, viinum, vota (pl), vulgus, vulnus.

REFERENCES

- ALLOPENNA, PAUL D.; JAMES S. MAGNUSON; and MICHAEL K. TANENHAUS. 1998. Tracking the time course of spoken word recognition using eye movements: Evidence for continuous mapping models. *Journal of Memory and Language* 38.419–39.
- ANDERSEN, ROGER. 1984. What's gender good for anyway? Second languages: A cross-linguistic perspective, ed. by Roger Andersen, 77–100. Rowley, MA: Newbury House.
- AYRES-BENNETT, WENDY. 1996. *A history of the French language through texts*. London: Routledge.
- BALDI, PHILIP. 1999. *The foundations of Latin*. Hawthorne: Mouton de Gruyter.
- BATALI, JOHN. 1998. Computational simulations of the emergence of grammar. Approaches to the evolution of language: Social and cognitive bases, ed. by James R. Hurford, Michael Studdert-Kennedy, and Chris Knight, 405–26. Cambridge: Cambridge University Press.
- BERMAN, RUTH. 1985. The acquisition of Hebrew. In Slobin 1985, 255–371.
- BIDOT, E. 1925. *La clef du genre des substantifs français. Méthode dispensant d'avoir recours au dictionnaire*. Poitiers: Impr. Nouvelle.
- BRIGHTON, HENRY. 2002. Compositional syntax from cultural transmission. *Artificial Life* 8.25–54.
- BYBEE, JOAN. 1985. *Morphology: A study of the relation between meaning and form*. Amsterdam: John Benjamins.
- BYBEE, JOAN. 1995. Regular morphology and the lexicon. *Language and Cognitive Processes* 10.425–55.
- CARROLL, SUSANNE E. 1995. The hidden dangers of computer modeling: Remarks on Sokolik and Smith's connectionist learning model of French gender. *Second Language Research* 11.193–205.
- CHOMSKY, NOAM, and MORRIS HALLE. 1968. *The sound pattern of English*. New York: Harper & Row.
- COHEN, MARCEL. 1973. *Histoire d'une langue: Le français*. Paris: Éditions sociales.
- CORBETT, GREVILLE. 1982. Gender in Russian: An account of gender specification and its relationship to declension. *Russian Linguistics* 6.197–232.
- CORBETT, GREVILLE. 1991. *Gender*. Cambridge: Cambridge University Press.
- CRAIG, COLETTE (ed.) 1986. *Noun classes and categorization*. Amsterdam: John Benjamins.
- DE GOROG, RALPH. 1982. *Dictionnaire inverse de l'ancien français*. Binghamton, NY: Center for Medieval & Early Renaissance Studies.
- DEMUTH, KATHERINE; NICOLAS FARACLAS; and LYNELL MARCHESE. 1986. Niger-Congo class and agreement systems in language acquisition and historical change. In Craig, 453–72.
- DIEDERICH, PAUL B. 1939. *The frequency of Latin words and their endings*. Chicago: University of Chicago Press.
- DOTTIN, GEORGES. 1920. *La langue gauloise: Grammaire, texts et glossaire*. Paris: Klincksieck.
- ELCOCK, W. D. 1960. *The Romance languages*. London: Faber & Faber.
- ELMAN, JEFFREY L. 1990. Finding structure in time. *Cognitive Science* 14.179–211.

- ELMAN, JEFFREY L. 1992. Grammatical structure and distributed representations. *Connectionism: Theory and practice*, ed. by Steven Davis, 138–78. New York: Oxford University Press.
- ERNOUT, ALFRED. 1927. *Morphologie historique du latin*. Paris: Klincksieck.
- ERNOUT, ALFRED. 1953. *Syntaxe latine*. Paris: Klincksieck.
- ERNOUT, ALFRED, and ANTOINE MEILLET. 1939. *Dictionnaire étymologique de la langue latine, histoire des mots*. Paris: Klincksieck.
- FLETCHER, P. C.; J. M. ANDERSON; D. R. SHANKS; R. HONEY; T. A. CARPENTER; T. DONOVAN; N. PAPADAKIS; and E. T. BULLMORE. 2001. Responses of human frontal cortex to surprising events are predicted by formal associative learning theory. *Nature Neuroscience* 4.1043–48.
- GASSER, MICHAEL. 1993. Learning words in time: Towards a modular connectionist account of the acquisition of receptive morphology. Indiana, IN: Computer Science TR–384.
- GODEFROY, FREDERIC. 1901. *Lexique de l'ancien français*. Paris and Leipzig: H. Welter.
- GROSJEAN, FRANÇOIS. 2000. The bilingual's language modes. One mind, two languages: *Bilingual language processing*, ed. by Janet Nicol, 31–55. Oxford: Blackwell.
- GUILLEMON, DELPHINE, and FRANÇOIS GROSJEAN. 2000. The gender effect in spoken word recognition: The case of bilinguals. Neuchâtel: University of Neuchâtel, ms.
- HARE, MARY, and JEFFREY L. ELMAN. 1995. Learning and morphological change. *Cognition* 56.61–98.
- HÄRMÄ, JUHANI. 2000. Gender in French: A diachronic perspective. *Gender in grammar and cognition*, ed. by Barbara Unterbeck, Matti Rissanen, Terttu Nervalainen, and Mirja Saari, 609–19. Berlin: Mouton de Gruyter.
- HARVEY, MARK, and NICHOLAS REID (eds.) 1997. *Nominal classification in Aboriginal Australia*. Amsterdam: John Benjamins.
- HOCKETT, CHARLES F. 1958. *A course in modern linguistics*. New York: Macmillan.
- JACKSON, DAN, and GARRISON COTTRELL. 1997. Attention and U-shaped learning in the acquisition of the past tense. *Proceedings of the Nineteenth Cognitive Science Conference*, 325–30. Mahwah, NJ: Lawrence Erlbaum.
- JOSEPH, BRIAN. 1979. On the animate-inanimate distinction in Cree. *Anthropological Linguistics* 21.351–54.
- JUSCZYK, PETER W.; ANGELA FRIEDERICI; JEANINE WESSELS; VIGDIS Y. SVENKERUD; and ANN-MARIE JUSCZYK. 1993. Infants' sensitivity to the sound patterns of native language words. *Journal of Memory and Language* 32.402–20.
- JUSCZYK, PETER W.; PAUL A. LUCE; and JAN CHARLES-LUCE. 1994. Infants' sensitivity to phonotactic patterns in native language. *Journal of Memory and Language* 33.630–45.
- KARMILOFF-SMITH, ANNETTE. 1979. *A functional approach to child language: A study of determiners and reference*. Cambridge: Cambridge University Press.
- KIRBY, SIMON. 1999. *Function, selection, and innateness: The emergence of language universals*. Oxford: Oxford University Press.
- KIRBY, SIMON. 2002. Natural language from Artificial Life. *Artificial Life* 8.185–215.
- LAKOFF, GEORGE. 1987. *Women, fire, and dangerous things*. Chicago: University of Chicago Press.
- LAMBERT, PIERRE-YVES. 1994. *La langue gauloise*. Paris: Éditions Errance.
- LAPOINTE, STEVEN G. 1988. Towards a unified theory of agreement. *Agreement in natural language*, ed. by Michael Barlow and Charles Ferguson, 67–87. Stanford: CSLI Publications.
- LEJEUNE, MICHEL. 1985. *Recueil des inscriptions gauloises*. T. 1: Textes gallo-grecs. Paris: Éditions du Centre national de la recherche scientifique.
- LEJEUNE, MICHEL. 1988. *Recueil des inscriptions gauloises*. T. 2: Textes gallo-étrusques, textes gallo-latins sur pierre. Paris: Éditions du Centre national de la recherche scientifique.
- LEJEUNE, MICHEL; EDOUARD BACHELLERY; FRANÇOISE BADER; and COLETTE BEMONT. 1985. Textes gaulois et gallo-romains en cursive latine, III: Le Plomb du Larzac. *Études celtiques* 22.95–177.
- LEWIS, CHARLTON T., and CHARLES SHORT. 1962. *A Latin dictionary*. Oxford: Clarendon Press.
- LING, CHARLES X. 1994. Learning the past tense of English verbs: The symbolic pattern associator vs. connectionist models. *Journal of Artificial Intelligence* 1.209–29.

- MACWHINNEY, BRIAN, and JARED LEINBACH. 1991. Implementations are not conceptualizations: Revising the verb learning model. *Cognition* 40.121–57.
- MALKIEL, YAKOV. 1975. Gender, sex, and size, as reflected in the Romance languages. *Studies in Romance linguistics*, ed. by Michio P. Hagiwara, 254–77. Rowley, MA: Newbury House.
- MANDLER, JEAN. 2000. Perceptual and conceptual processes in infancy. *Journal of Cognition and Development* 1.3–36.
- MARCUS, GARY F.; MICHAEL ULLMAN; STEVEN PINKER; MICHELLE HOLLANDER; T. JOHN ROSEN; and FEI XU. 1992. Overregularization in language acquisition. *Monographs of the Society for Research in Child Development* 57.4.
- McKITTERICK, ROSAMOND. 1989. *The Carolingians and the written word*. Cambridge: Cambridge University Press.
- MEILLET, ANTOINE. 1966. *Esquisse d'une histoire de la langue latine*. Paris: Klincksieck.
- MEL'ČUK, IGOR A. 1974. Statistics and the relationship between the gender of French nouns and their endings. *Essays on lexical semantics 1*, ed. by Victor Ju. Rozenčevjg, 11–42. Stockholm: Skriptor.
- MEYER-LUBKE, WILHELM. 1923. *Grammaire des langues romanes, t. 2: Morphologie*. New York: G. E. Stechert.
- MODER, CAROL. 1992. *Productivity and categorization in morphological classes*. Buffalo, NY: SUNY dissertation.
- MYLNE, TOM. 1995. Grammatical category and world view: Western colonization of the Dyrbal language. *Cognitive Linguistics* 6.379–404.
- NICHOLS, JOHANNA. 1989. The Nakh evidence for the history of gender in Nakh-Daghestanian. *The non-Slavic languages of the USSR: Linguistic studies NSL-5*, ed. by Howard I. Aronson, 158–75. Chicago: Chicago Linguistic Society.
- PEDERSEN, HOLGER. 1913. *Vergleichende Grammatik der keltischen Sprachen*. Bd. 2. Göttingen: Vandenhoeck und Ruprecht.
- PINKSTER, HARM. 1990. *Latin syntax and semantics*. London: Routledge.
- PLAUT, DAVID C. 1999. Computational modeling of word reading, acquired dyslexia, and remediation. *Converging methods in reading and dyslexia*, ed. by Raymond M. Klein and Patricia A. McMullen, 339–72. Cambridge, MA: MIT Press.
- PLUNKETT, KIM, and VIRGINIA MARCHMAN. 1991. U-shaped learning and frequency effects in a multilayered perceptron: Implications for child language acquisition. *Cognition* 38.43–102.
- PLUNKETT, KIM, and VIRGINIA MARCHMAN. 1993. From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition* 48.21–69.
- POLINSKY, MARIA, and DAN JACKSON. 1999. Noun classes: Language change and learning. *Cognition and function in language*, ed. by Barbara Fox, Dan Jurafsky, and Laura Michaelis, 29–50. Stanford: CSLI Publications.
- POPE, MILDRED K. 1961. *From Latin to Modern French*. Manchester: University of Manchester Press.
- ROHLFS, GERHARD. 1960. *Vom vulgärlatein zum altfranzösischen*. Tübingen: Max Niemeyer.
- RUMELHART, DAVID E.; JAMES L. McCLELLAND; and THE PDP RESEARCH GROUP (eds.) 1986. *Parallel distributed processing: Explorations in the microstructure of cognition*. Vol. 1: Foundations. Cambridge, MA: MIT Press.
- SAFFRAN, JENNY R.; RICHARD N. ASLIN; and ELISSA L. NEWPORT. 1996. Statistical learning by 8-month-old infants. *Science* 274.1926–28.
- SAFFRAN, JENNY R.; ELISSA L. NEWPORT; and RICHARD N. ASLIN. 1996. Word segmentation: The role of distributional cues. *Journal of Memory and Language* 35.606–21.
- SANKOFF, GILLIAN. 1980. *The social life of language*. Philadelphia: University of Pennsylvania Press.
- SCHMIDT, K.-H. 1967. Keltisches Wortgut im lateinischen. *Glotta* 44.151–74.
- SLOBIN, DAN I. 1973. Cognitive prerequisites for the development of grammar. *Studies of child language development*, ed. by Charles Ferguson and Dan I. Slobin, 175–208. New York: Holt, Rinehart & Winston.
- SLOBIN, DAN I. (ed.) 1985. *The cross-linguistic study of language acquisition*. Hillsdale, NJ: Lawrence Erlbaum.
- SMITH, PHILIP T. 1995. Are morphemes really necessary? Morphological aspects of language processing, ed. by Laurie Beth Feldman, 365–82. Hillsdale, NJ: Lawrence Erlbaum.

- SMOCZYŃSKA, MAGDALENA. 1985. The acquisition of Polish. In Slobin 1985, 595–685.
- SOKOLIK, M. E., and MICHAEL E. SMITH. 1992. Assignment of gender to French nouns in primary and secondary language: A connectionist model. *Second Language Research* 8.39–58.
- TAESCHNER, TRAUDE. 1983. *The sun is feminine: A study on language acquisition in bilingual children*. Berlin: Springer-Verlag.
- THOMASON, SARAH G., and TERRENCE KAUFMAN. 1988. *Language contact, creolization, and genetic linguistics*. Berkeley: University of California Press.
- TUCKER, G. RICHARD; WALLACE E. LAMBERT; and ANDRÉ A. RIGAULT. 1977. The French speaker's skill with grammatical gender: An example of rule-governed behavior. The Hague: Mouton.
- VÄÄNÄNEN, VEIKKO. 1981. *Introduction au latin vulgaire*. Paris: Klincksieck.
- VAN BERKUM, JOS J. A.; COLIN M. BROWN; and PETER HAGOORT. 1999. When does gender constrain parsing? Evidence from ERPs. *Journal of Psycholinguistic Research* 28.555–71.
- VAN EVERBROECK, EZRA. 2003. Language type frequency and learnability from a connectionist perspective. *Linguistic Typology* 7.1–50.
- VINCENT, NIGEL. 1988. Latin. *The Romance languages*, ed. by Martin Harris and Nigel Vincent, 26–78. London: Croom Helm.
- WAGNER, ROBERT LEON. 1974. *L'Ancien français: points de vue, programmes*. Paris: Larousse.
- WALSH, MICHAEL. 1993. Classifying the world in an Aboriginal language. *Language and culture in Aboriginal Australia*, ed. by Michael Walsh and Colin Yallop, 107–22. Canberra: Aboriginal Studies Press.
- WARTBURG, WALTHER VON. 1888 [1962]. *Evolution et structure de la langue française*. 6. edn. Berne: A. Francke.
- WHATMOUGH, JOSHUA. 1963. *The dialects of Ancient Gaul: Grammar*, vol. 1. Ann Arbor: Edwards Brothers.
- WHATMOUGH, JOSHUA. 1970. *The dialects of Ancient Gaul*. Cambridge, MA: Harvard University Press.
- WICHA, NICOLE Y. Y. 2002. *Grammatical gender in real-time language comprehension in Spanish: Behavioral and electrophysiological investigations*. La Jolla, CA: University of California, San Diego, dissertation.
- WICHA, NICOLE Y. Y.; EVA M. MORENO; and MARTA KUTAS. 2003. Expecting gender: Event related brain potentials to the expectation of a specific gender marked article in a Spanish sentence. *Cortex* 39, to appear.
- WRIGHT, ROGER. 1982. *Late Latin and Early Romance in Spain and Carolingian France*. Liverpool: Francis Cairns.
- ZUBIN, DAVID A., and KLAUS-MICHAEL KÖPCKE. 1986. Gender and folk taxonomy: The indexical relation between grammatical and lexical categorization. In Craig, 139–80.

Department of Linguistics, 0108
 University of California San Diego
 9500 Gilman Drive
 La Jolla, CA 92093-0108
 [polinsky@ling.ucsd.edu]
 [ezra@ling.ucsd.edu]

[Received 19 March 2001;
 revision received 2 July 2002;
 accepted 6 December 2002]