Inducing Concatenative Units from Machine Readable Dictionaries and Corpora for Speech Synthesis

Evelyne Tzoukermann* and Judith L. Klavans†*

†Columbia University, Department of Computer Science
New York, New York 10027
klavans@cs.columbia.edu
and
*A.T.&T. Bell Laboratories
600 Mountain Avenue, Murray Hill, N.J. 07974
klavans@research.att.com
evelyne@research.att.com

Abstract

The purpose of this research is to determine the best method for deciding on an optimal set of concatenative units for concatenative speech synthesis. Of the two main approaches to speech synthesis: segmental synthesis and rule-based synthesis, the former relies heavily on the successful choice of concatenative units. Segmental synthesis consists of concatenating segmental units (diphones, triphones, etc); rule-based synthesis consists of the computation of control parameters based on pre-established rules.

Deciding on the set of diphones is quite straightforward in the sense that it suffices to take the phoneme inventory of a language, and simply combine each phoneme with every other one. For example, taking the approximately 35 French phonemes, 1225 phonemic pairs (35x35) constitute the complete and exhaustive starting diphone inventory. On the other hand, deciding on the set of triphones, quadriphones and larger units raises difficult questions about the nature of phonemes in a given language such as: (1) stability vs instability in a coarticulatory environment, (2) size of overall inventory, and (3) frequency of that unit in the language, in combination with factors (1) and (2).

We report on experiments with four different databases, with comparisons between the resources regarding their n-gram frequency output. The first two databases consist of pronunciation field information from two dictionaries, the Encyclopedic Robert French dictionary [16] with 85,000 headwords, and the smaller Collins Gem [13] containing 15,000 words. For comparison, we use two text corpora, the Hansard (about 2.5 million words) and the smaller Tubach and Boe [31] corpus (80,000 words); both corpora were processed by a set of grapheme-to-phoneme rules [18]. A frequency extraction program was applied to all four resources to extract trigram phonemic frequencies; this serves as a basis for comparison between dictionary derived data and corpus derived frequencies.

Introduction

In this paper, we explore methods to capture an efficient set of concatenative units, given several sources of input, so that it can be incorporated in the French NewTTS system that is in the process of being developed at AT&T Bell Laboratories. Our purpose is two-fold: one practical and one theoretical. On the practical side, a speech synthesis system for French is being built at AT&T Bell Laboratories [17], and the output of this research will help to determine and optimize set of most frequent concatenative units for putting into the working system. On the theoretical side, our goal is to explore the question of what constitutes adequate data for inducing this optimal set of concatenative units. Once the frequency program is run on all the data, we compare these different output on several dimensions:

1. A comparison of the triphone frequencies run over dictionary pronunciations from a small and large dictionary;
2. A comparison of the triphone frequencies run over the Hansard corpus and the Tubach and Boe corpus;
3. A comparison of the triphone frequencies resulting from the output of 1, i.e. the total dictionary data, and output of 2, i.e. the total corpus data.

Point 1 is particulary important since there is some doubt about the advantage gained from using dictionary pronunciations. If grapheme-to-phoneme rules provide the same results as pronunciation fields themselves, then the justification for using both resources is greatly reduced since once would be adequate. Traditionally, it is the MRD that is often criticized but the corpus would be equally dispensable. One should not neglect the fact that MRDs are presently more available than complete grapheme-to-phoneme systems, and that alone justifies their use either as a first pass set of pronunciations or as data for rule generation and testing. In order to generalize our conclusions, of course a full comparison of dictionary pronunciations and grapheme-to-phoneme output would have to be done over several languages, with different phonemic and stress properties, since one language alone is not sufficient to determine the larger question on the utility of MRDs.

Added to the question of the use of MRD data is the issue of the role of corpus data in text-to-speech (TTS). Typically, TTS systems are rule-based, the most common set of steps being: morphological analysis of word forms by rule, word level pronunciation by grapheme-to-phoneme rules, interword level phenomena by rule, and intonational contouring, also by rule. Corpus data has just started being used for developing TTS systems [1] and we report on one of these usages. Although the majority of speech recognition systems rely on probabilistic methods over large text corpora, such techniques have been slow to find their ways into TTS systems. This paper demonstrates one specific way that such methods are useful in combination with a rule-based system. Thus, we present here a novel combination approach, one which uses both probabilistic and rule-based approaches for building a TTS system.

The Use of Machine-Readable Dictionaries

MRDs, no matter how small and no matter how internally inconsistent, have proven to be an invaluable resource for a wide variety of applications [2]. Machine readable dictionaries first became available on a limited basis in the 60's. Merriam Webster's Seventh (W7) [25] was manually input and was informally passed around on magnetic tape and in long print-outs from site to site. Other dictionaries (e.g. LDOCE, Collins Bilinguals, Robert monolingual, and so on) continue to become available for use.

To date, TTS research has primarily used Webster's Seventh (for English), Collins' bilinguals, and other non-English monolinguals, since these are available in electronic format.[1] The diphone based text to speech system first developed at the IBM T. J. Watson research Center in the late 60's [15] later added syntactic and morphological information from W7. The later system, described

[1] Mention must be made of the CoBuild Dictionary [9] which promises to be the most useful so far, but due to its expense, has turned out to be of limited use at this time, except by large very well-funded industrial laboratories (e.g. Microsoft, AT&T Bell Laboratories).

in [30] was sentence-based; each sentence was run through a syntactic parser for part of speech assignment. Pronunciations were retrieved from a list of spellings and pronunciations derived and cleaned from W7. If a form could not be found, first a check was made to see if it was a punctuation mark, number, etc.; these were handled by special routines. If the spelling was none of the above, it was pronounced by the spelling-to-phoneme routines [10], [21]. Pre-processing took care of acronyms and abbreviations, and some post-processing was done (e.g., to flap t's and d's where appropriate and where not indicated in W7). [8] reports on the successful use of W7 for a stress assignment guessing algorithm, as part of the TTS system at AT&T Bell Laboratories. Carter [5] discusses the use of LDOCE pronunciation field information for speech recognition. The next section of the paper focuses on the structure of electronic dictionaries for TTS in particular.

Using MRDs in Text to Speech: Since a dictionary entry consists of several fields of information, naturally, each will be useful for different applications. Among the standard fields are pronunciation, etymology, subject field notes, definition fields, synonym and antonym cross references, semantic and syntactic comments, run-on forms, conjugational class and inflectional information where relevant, and translation for the bilingual dictionaries. The most directly useful for TTS is the pronunciation field. Equally useful, but less directly accessible, are data from run-on fields, conjugational class information, and part-of-speech. [2]

To illustrate, the following partial entries from W7 illustrate typical pronunciation, definition, and run-on fields:

(1) **ha.ven** /ˈh.ā-vən/ *n* **1**: HARBOR, PORT **2** : a place of safety : ASYLUM – **haven** *vt*
(2) **den.tic.u.late** /den-ˈtik-yə-lət/ *or* **den.tic.u.lat.ed** /-ˌlāt-əd/ *adj*

The entry for "haven" contains one full pronunciation, with several definitions, some synonyms, and two parts of speech. The entry for "denticulate" has a morphologically related run-on form "denticulated" in the early part of the entry, and the pronunciation of that run-on is related to the main entry, but the user must figure out how to chop and append the given syllables. [7] reports on the need to resyllabify entries already syllabified in LDOCE, since syllable boundaries for written forms usually reflect hyphenation conventions, rather than phonologically motivated syllabification conventions necessary for pronunciation. While these types of reasoning are not difficult for humans, for whom the dictionary was written, they are quite difficult for programs, and thus are not straighforward to perform automatically.

The pronunciation field: Extracting the pronunciation field from the MRD is one of the most obvious uses of the dictionary. Nevertheless, parsing dictionaries in general can be a very complex operation [26] and even the extraction of one field, such as the pronunciation, can pose problems. Similar to W7, in the Robert French dictionary [16], which contains about 89,000 entries, several pronunciations can be given for a headword and the choice of one must be made. Moreover, because of the rich morphology of French which has a rough ratio of eight morphologically inflected words for one baseform, Robert lists only the non-inflected forms of the lexical entries. However, if pronunciation varies during inflection of nouns and adjectives, the pronunciation field reflects that variation which makes the information difficult to extract automatically. For example, in (3) and (4), one needs to know the nature of the rule to apply in order to get the feminine form of the adjective.

(3) **blanc, blanche** /blɑ̃, blɑ̃ʃ/ adj. et n.
(4) **vif, vive** /vif, viv/ adj. et n.

In (3), the feminine /blɑ̃ʃ/ (blanche,"white") is obtained by adding the phoneme /ʃ/ to the masculine /blɑ̃/. In (4), the form /viv/ (vive,"sharp, quick") is formed by stripping the affix /f/ and adding the morpheme /v/. For the French TTS system [17], the set of diphones was established by taking most of the phonemes for French (35) and coupling them with another ($35^2 = 1225$ pairs). Then, the diphones were extracted from the pronunciation field of the headword of the Robert dictionary. Out of 1225 phonemic pairs, 874 words were found with at least one occurence of the pair. This number is explained by the fact that the dictionary contains (a) only non-inflected words, and (b) lists of words as opposed to words in texts, therefore no inter-word triphones.

A program was written to search through the dictionary phonetic field and select the longest word where the phoneme pairs would be in middle syllable position. The pair [headword_orth, headword_phon] was extracted and headword_orth was placed in a carrier sentence for recording. Due to the lack of explicit inflectional information for nouns and adjectives, only the non-inflected forms of the entries were extracted during dictionary lookup for building the diphone table. Similarly for verbs, only the infinitive forms were used since the dictionary does not list the inflected forms as headwords. This exemplifies the most simple way to use pronunciation field data, which we have completed. A pronunciation list of around 85,796 phonetic words was obtained from the original list of almost 89,000 entries, i.e. 96% of the entries. The remaining 4% consist primarily of prefixes and suffixes which are listed in the dictionary without pronunciations and which would not be useful in isolation in any case.

Related Research

Dictionary definitions have been used as the input to statistical methods for word sense disambiguation by [32], among others. Definitions have also been used as input to probabilistic systems [6], and as a source of relational data [23] for building lexical knowledge bases. The question of the relationship between dictionary data and the larger language is raised by [22] in building bilingual lexicons, again focussing on definition and translation fields. Although the statistical analysis of MRDs has focussed primarily on definitions and translations, [19] has used the pronunciation field as data.

Strategies for finding a methodology for building polyphonic units differ greatly. Some researchers, such as [3] based their set on perceptually validated tests; the perceptually salient categories were determined and the final set was selected using response frequencies grouped by categories. Others, such as [4], first group the different phonemes by classes, separating them according to their acoustic stability or unstability. Each phoneme is combined with another and then the unstable ones are surrounded by stable phonemes. In the end, a collection of over 7700 polyphones is collected and statistics are run through a corpus to select the top 3000 units.

In [19], a dictionary of over 110,000 entries containing 51,219 common words and 59,625 proper nouns is used for selecting candidate units [27], [11], [28]. The phonemic string is split according to ten language-dependent segmentation principles. For example, the word "abacus" [ˈab-ə-kəs] is first transformed into cuttable units as follows: [#ˈa,ˈab,bə,ək,kə,əs,s#]. Once each dictionary word is split, the duplicates are removed and the remaining units form the set of concatenative units. At the end of this operation, a rather long list is obtained that can be pruned by methods such as acoustic clustering, reduction of secondary and primary stress into one stress in order to keep only one +stress/-stress dictinc-

[2] Notice, however, that the full Collins Spanish-English dictionary [12], as opposed to the other bilinguals, does not contain any pronunciation information. Although this is rather surprising taking into account that the smaller versions such as the paperback and gem ([14], [13], [24]) do have a phonetic field, it could be attributed to the fact that pronunciation rules in Spanish are relatively predictable.

tion. Other techniques allow the selection of a minimal set of word pairs for inter-word junctures; every candidate unit inside and across words is included. The same strategy was replicated on the Collins Spanish-English dictionary by [20]. In this fashion, the dictionary is used as a sample of the language in the sense that it assumes that most of the phonemic combinations in the language are present.

Experiments

Our experiments consisted of taking four different databases (two French dictionaries and two French corpora) and running a frequency extraction program on each of these. First, each database was processed and carefully cleaned to remove unwanted data; then, grapheme-to-phoneme conversion software was run so that the database output would be comparable.

The different output appear as a sorted list by triphone frequency of sentences or phrases, with triphones occuring in each phrase. Out of this list, we kept the top 1,000 triphones. Each corpus was processed in the same fashion; the triphones occured as follows:

Hansard corpus 2.5 million w.		**Tubach Boe corpus 80,000 words**		**Robert dict 85,000 words**		**Collins dict 15,000 words**	
44956	sjO	1193	iks	3636	mA*	1096	mA*
28582	jO*	1151	sEt	3324	ik*	1093	te*
28097	par	904	sjO	3223	jO*	1078	jO*
26163	asj	723	avE	2823	sjO	1013	*de
25789	pr>	687	d&l	2597	te*	1012	sjO
24094	tr&	655	ilj	2202	*de	811	je*
20555	&la	651	&la	2105	Ysj	794	@r*
20472	pur	632	bjI	2086	Er*	791	*kO
18109	&mA	630	k>m	2067	aZ*	694	Asj
17713	d&l	619	sk&	1789	ist	652	Er*
16972	*Z&	610	al>	1725	ism	551	ik*
16760	sEt	606	etE	1554	ite	544	re*
16717	k>m	596	vwa	1492	je*	509	*ka
15900	*mE	593	pas	1462	sm*	489	*pr
15841	te*	574	ply	1405	bl*	482	ne*
15321	&pr	558	lja	1391	ali	477	bl*
14939	ist	504	ask	1389	abl	447	le*
14865	*s&	500	kil	1376	tik	447	ite
14623	Est	498	mEm	1341	*ka	443	ist
14523	mA*	496	AfI	1321	st*	431	ze*
14313	*sE	492	asj	1301	tri	412	abl
13811	Er*	492	EpY	1262	@r*	410	*ma
13714	*il	462	ite	1186	ri*	402	ir*
13712	ite	455	ala	1147	*pr	380	ar*
13462	pre	445	mwa	1084	lit	375	al*

Table 1: Twenty five most frequent triphones

Table 1 shows the top 25 triphones of the four databases, each triphone being associated with its number of occurences in the corpus. The Robert and Collins dictionaries contain the same top triphone "mA*" (-ment) probably due to a proportionally similar number of adverbs (since most of French adverbs end in "-ment"). It is interesting to point out that the triphone "sjO" ("tion" in French) occurs in the four databases among the top 5 triphones.

Results

Considering the widely different distributions of top triphones, we decided to measure the level of commonality between the different sets. The level of commonality was defined as a percentage corresponding to the triphones in common in two different sets of the same size. We split the lists of the most frequent triphones from each database into smaller set of 10, 100, and 1000 triphones. Each set was compared to another set of the same size. For example, the ten top triphones from the Robert dictionary were compared to the ten top triphones of the Collins, the Tubach and Boe, and the Hansard. Additionally, in order to obtain a more informative comparison of text type, we classed the two dictionary sets into one and the two database sets into one to compare the two dictionaries with the two databases. Interestingly, as shown in Tables 2 and 3, the level of commonality between all sets never exceeds 35%. Nevertheless, it is interesting to point out that the highest level of commonality occurs between the two dictionaries (from 30% to 35%). It is understandable indeed that in comparing two dictionaries of two very different sizes, the words of the smaller may be a balanced subset of the larger. However, it is surprising that there was no more than 35% overlap. The two databases have a fair number of triphones in common, even though the top ten triphones have 25% overlap, the top 100 triphones dip to only 13% overlap, and the top 1000 triphones have only 29% overlap. The comparison between the two databases and the two dictionaries did not appear to be very informative. It was rather comparable to the other databases.

		Collins 25K			**Tub/Boe 15K**			**Hansard 2.5M**		
		10	100	1000	10	100	1000	10	100	1000
Robert 85k	10	30%	–	–	27.5%	–	–	05%	–	–
	100	–	35%	–	–	04%	–	–	05%	–
	1000	–	–	35%	–	–	16%	–	–	15%
Collins 25K	10	–	–	–	05%	–	–	10%	–	–
	100	–	–	–	–	06%	–	–	08%	–
	1000	–	–	–	–	–	15%	–	–	16%
Tub/Boe 15K	10	–	–	–	–	–	–	25%	–	–
	100	–	–	–	–	–	–	–	13%	–
	1000	–	–	–	–	–	–	–	–	29%

Table 2: Measure of commonality of the most frequent trigrams between the four databases (one by one comparison)

		Rob-Coll			**Tub/Boe-Hans**		
		10	100	1000	10	100	1000
Robert 85K	10	–	–	–	17%	–	–
	100	–	–	–	–	13%	–
	1000	–	–	–	–	–	31%
Collins 25K	10	–	–	–	17%	–	–
	100	–	–	–	–	16%	–
	1000	–	–	–	–	–	31%
Tub/Boe 15K	10	24%	–	–	–	–	–
	100	–	27%	–	–	–	–
	1000	–	–	35%	–	–	–
Hansard 2.5M	10	27%	–	–	–	–	–
	100	–	29%	–	–	–	–
	1000	–	–	35%	–	–	–
Rob-Coll	10	–	–	–	27.5%	–	–
	100	–	–	–	–	30%	–
	1000	–	–	–	–	–	43%

Table 3: Measure of commonality of the most frequent trigrams between the four databases and the two dictionaries and the two corpora

Final Remarks

Unlike the word set in a language, the set of phonemes/phonetic units in a language is significantly smaller. Furthermore, there is no morphological productivity to create the effects of flocking or to create new units [29]. Thus, it was our hypothesis that a dictionary could provide essential concatenative unit data without recourse to larger corpora. Our results showed that for interword phenomena, the hypothesis was correct, whereas for intraword effects, such as liaison and assimilation, the dictionary field information falls short. To make a more accurate count, it would be necessary to compute the number of triphones found in the corpora that come solely from inter-word boundaries.

Our explanation for the apparently small overlap relates to the size of the selected sets (10, 100, 1000 to triphones) as related to the possible triphone unit set. The set of French triphones (35^3 = 42,875 units) is a large number and it is not surprising if the top 1000 (e.g. 0.02% of 42,875) come out to be different in each database. Testing the top 20,000 triphones would surely gather more commonality. Nevertheless, for practical purposes, such a number would be difficult to handle in a TTS system.

Our next steps include determining methods for evaluating choices between competing selected sets. In addition we intend to consider linguistic phenomena such as transitions between unstable phonemes (such as liquids and glides). We believe that approaches, such as [4] that use linguistic knowledge and select the output using statistics on corpora, can help optimization. However, we also believe that building an optimal set of units is an iterative process that is a function of the synthesized voice itself as well as the synthesis technique (e.g. TD-PSOLA or LPC) and must be properly evaluated.

References

[1] Benjamin Ao, Chilin Shih, and Richard Sproat. A corpus-based mandarin text-to-speech synthesizer. Yokohama, Japan, 1994. Proceedings of the 2nd International Conference on Spoken Language Processing.

[2] Beryl T. Atkins, Judith Kegl, and Beth Levin. Anatomy of a verb entry: from linguistics theory to lexicographic practice. *International Journal of Lexicography*, 1:84–126, 1988.

[3] Veronique Aubergé. *La synthèse de la parole: 'des règles aux lexique'.* PhD thesis, Université de Grenoble, Grenoble, France, 1991.

[4] Frédéric Bimbot. *Synthèse de la Parole: Des Segments aux règles, avec utilisation de la décomposition temporelle.* PhD thesis, Telecom Paris 88 E019, Paris, France, 1988.

[5] Bran Boguraev and Ted Briscoe. *Introduction*, chapter Chapter I, pages 1–40. Longman, Burnt Hill, Harlow, Essex, 1989.

[6] P. Brown, S. Della Pietra, V. Della Pietra, M. Goldsmith, J. Hajic, R. Mercer, and S. Mohanty. But dictionaries are data too. In *Proceedings of DARPA*, Princeton, New Jersey, 1993. Workshop on Speech and Natural Language.

[7] David Carter. *LDOCE and Speech Recognition*, chapter Chapter 6, pages 135–152. Longman, Burnt Hill, Harlow, Essex, 1989.

[8] Kenneth Church. Stress assignment in letter to sound rules for speech synthesis. In *ACL Proceedings, 23rd Annual Meeting*, pages 246–253, Morristown, NJ, 1985. Association for Computational Linguistics.

[9] *Collins Cobuild English Language Dictionary.* Collins Publishers, London, 1987.

[10] Paul Cohen. Spelling to sound conversion for text to speech. 1982.

[11] Cecil Coker. A dictionary-intensive letter-to-sound program. *Journal of the Acoustical Society of America, Supplement 1*, 78:S7, 1985.

[12] *Collins Spanish Dictionary: Spanish-English.* Collins Publishers, Glasgow, 1989.

[13] P-H Cousin. *Collins Gem French Dictionary: French-English. English-French.* Harper Collins Publishers, London, 1991.

[14] P-H. Cousin, L. Sinclair, J-F. Allain, and C. E. Love. *The Collins Paperback French Dictionary: French-English. English-French.* Collins Publishers, London, 1989.

[15] N.R. Dixon and H. D. Maxey. Terminal analog synthesis of continuous speech using the diphone method of segment assembly. In *IEEE Transactions on Audio and Electroacoustics AU16*, pages 40–50, 1968.

[16] Alain Duval et al. *Robert Encyclopedic Dictionary (CD-ROM).* Hachette, Paris, 1992.

[17] Tzoukermann Evelyne. Issues in text-to-speech for french. In *Proceedings of Coling94*, Kyoto, Japan, 1994. International Conference on Computational Linguistics.

[18] Marty F. Trois systèmes informatiques de transcription phonétique et graphémique. *Le Fran cais Moderne*, LX, 2:179–197, 1992.

[19] Coleman John. Computation of candidate synthesis units. In *11222-930719-07TM*, Murray Hill, N.J., USA, 1993. Technical Memorandum, AT& Bell Laboratories.

[20] Coleman John and Pilar Prieto. Accurate pronunciation rules for american spanish text-to-speech. In *11222-930719-06TM*, Murray Hill, N.J., USA, 1993. Technical Memorandum, AT& Bell Laboratories.

[21] Judith Klavans and Sara Basson. Documentation of letter to sound components of the walrus text to speech system. Yorktown Heights, New York, 1984. Unpublished manuscript, IBM, T. J. Watson Research Center.

[22] Judith Klavans and Evelyne Tzoukermann. The bicord system: Combining lexical information from bilingual corpora and machine readable dictionaries. Helsinki, Finland, 1990. Proceedings of the 13th International Conference on Computational Linguistics.

[23] Judith L. Klavans, Martin S. Chodorow, and Nina Wacholder. From dictionary to knowledge base via taxonomy. Centre for the New Oxford English Dictionary and Text Research: Electronic Text Research, University of Waterloo, Canada, 1990. Proceedings of the Sixth Conference of the University of Waterloo.

[24] Gonzales M. *Collins Gem Spanish Dictionary: French-English. English-French.* Harper Collins Publishers, London, 1990.

[25] Merriam. *Webster's Seventh New Collegiate Dictionary.* G.& C. Merriam, Springfield, Mass., 1963.

[26] M. Neff and B. Boguraev. Dictionaries, dictionary grammars and dictionary entry parsing. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, Vancouver, Canada, 1989. Association for Computational Linguistics.

[27] Olive J. P. A new algorithm for a concatenative speech synthesis system using an augmented acoustic inventory of speech sounds. In Gérard Bailly and Christian Benoit, editors, *Proceedings of the ESCA Workshop on Speech Synthesis*, 1990.

[28] Olive J. P. and M. Y. Liberman. A set of concatenative units for speech synthesis. In J. J. Wolf and D. H. Klatt, editors, *Speech Communication Papers Presented at the 97th Meeting of the Acoustical Society of America*, pages 515–518, New York: American Institute of Physics, 1979.

[29] Richard W. Sproat. Newtts: A user's manual. In *11222-921102-09TM*, Murray Hill, N.J., USA, 1992. Technical Memorandum, AT& Bell Laboratories.

[30] John Thomas, Judith Klavans, Jonas Nartey, Cliff Pickover, David Reich, and Mary Beth Rosson. "walrus: High-quality text-to-speech research system. pages 19–28. Proceedings of the IEEE on Speech Synthesis and Recognition, 1984.

[31] J-P Tubach and Boe L J. Un corpus de transcription phonétique (300,000 phones) - constitution et exploitation statistique. Ecole Nationale Supèrieure des Télécommunications, 1990. Telecom Paris 90 D 002.

[32] Yorick Wilks, Dan Fass, Cheg-Ming Guo, James McDonald, Tony Plate, and Brian Slator. *A tractable machine dictionary as a resource for computational semantics*, chapter Chapter 9, pages 193–228. Longman, Burnt Hill, Harlow, Essex, 1989.