

Chapter 8

Speech Synthesis

And computers are getting smarter all the time: Scientists tell us that soon they will be able to talk to us. (By 'they' I mean 'computers': I doubt scientists will ever be able to talk to us.)

Dave Barry

In Vienna in 1769, Wolfgang von Kempelen built for the Empress Maria Theresa the famous Mechanical Turk, a chess-playing automaton consisting of a wooden box filled with gears, behind which sat a robot mannequin who played chess by moving pieces with his mechanical arm. The Turk toured Europe and the Americas for decades, defeating Napoleon Bonaparte and even playing Charles Babbage. The Mechanical Turk might have been one of the early successes of artificial intelligence were it not for the fact that it was, alas, a hoax, powered by a human chess player hidden inside the box.

What is less well known is that von Kempelen, an extraordinarily prolific inventor, also built between 1769 and 1790 what was definitely not a hoax: the first full-sentence speech synthesizer. His device consisted of a bellows to simulate the lungs, a rubber mouthpiece and a nose aperture, a reed to simulate the vocal folds, various whistles for the fricatives, and a small auxiliary bellows to provide the puff of air for plosives. By moving levers with both hands to open and close apertures, and adjusting the flexible leather "vocal tract", an operator could produce different consonants and vowels.

More than two centuries later, we no longer build our synthesizers out of wood and leather, nor do we need human operators. The modern task of **speech synthesis**, also called **text-to-speech** or **TTS**, is to produce speech (acoustic waveforms) from text.

Modern speech synthesis has a wide variety of applications. Synthesizers are used in telephone-based conversational agents that conduct dialogues with people (Chapter 24). Synthesizers are also important in non-conversational applications that speak **to** people, such as in devices that read out loud for the blind or in video games or children's toys. Finally, speech synthesis can be used to speak **for** sufferers of neurological disorders, such as astrophysicist Steven Hawking who, having lost the use of his voice because of ALS, speaks by typing to a speech synthesizer and having the synthesizer speak the words. State of the art systems in speech synthesis can achieve remarkably natural speech for a very wide variety of input situations, although even the best systems still tend to sound wooden and are limited in the voices they use.

The task of speech synthesis is to map a text like the following:

(8.1) PG&E will file schedules on April 20.

to a waveform like the following:



Speech synthesis

Text-to-speech

TTS

Text analysis
Waveform synthesis

Speech synthesis systems perform this mapping in two steps, first converting the input text into a **phonemic internal representation** and then converting this internal representation into a waveform. We will call the first step **text analysis**, and the second step **waveform synthesis** (although other names are also used for these steps).

A sample of the internal representation for this sentence is shown in Fig. 8.1. Note that the acronym PG&E is expanded into the words P G AND E, the number 20 is expanded into *twentieth*, a phone sequence is given for each of the words, and there is also prosodic and phrasing information (the *'s), which we define later.

			*				*				*	L-L%																							
P	G	AND	E	WILL	FILE	SCHEDULES	ON	APRIL	TWENTIETH																										
p	iy	jh	iy	ae	n	d	iy	w	ih	l	f	ay	l	s	k	eh	jh	ax	l	z	aa	n	ey	p	r	ih	l	t	w	eh	n	t	iy	ax	th

Figure 8.1 Intermediate output for a unit selection synthesizer for the sentence *PG&E will file schedules on April 20.* The numbers and acronyms have been expanded, words have been converted into phones, and prosodic features have been assigned.

While text analysis algorithms are relatively standard, there are three widely different paradigms for waveform synthesis: **concatenative synthesis**, **formant synthesis**, and **articulatory synthesis**. The architecture of most modern commercial TTS systems is based on concatenative synthesis, in which samples of speech are chopped up, stored in a database, and combined and reconfigured to create new sentences. Thus, we focus on concatenative synthesis for most of this chapter, although we briefly introduce formant and articulatory synthesis at the end of the chapter.

Hourglass metaphor

Figure 8.2 shows the TTS architecture for concatenative unit selection synthesis, using the two-step **hourglass metaphor** of Taylor (2008). In the following sections, we examine each of the components in this architecture.

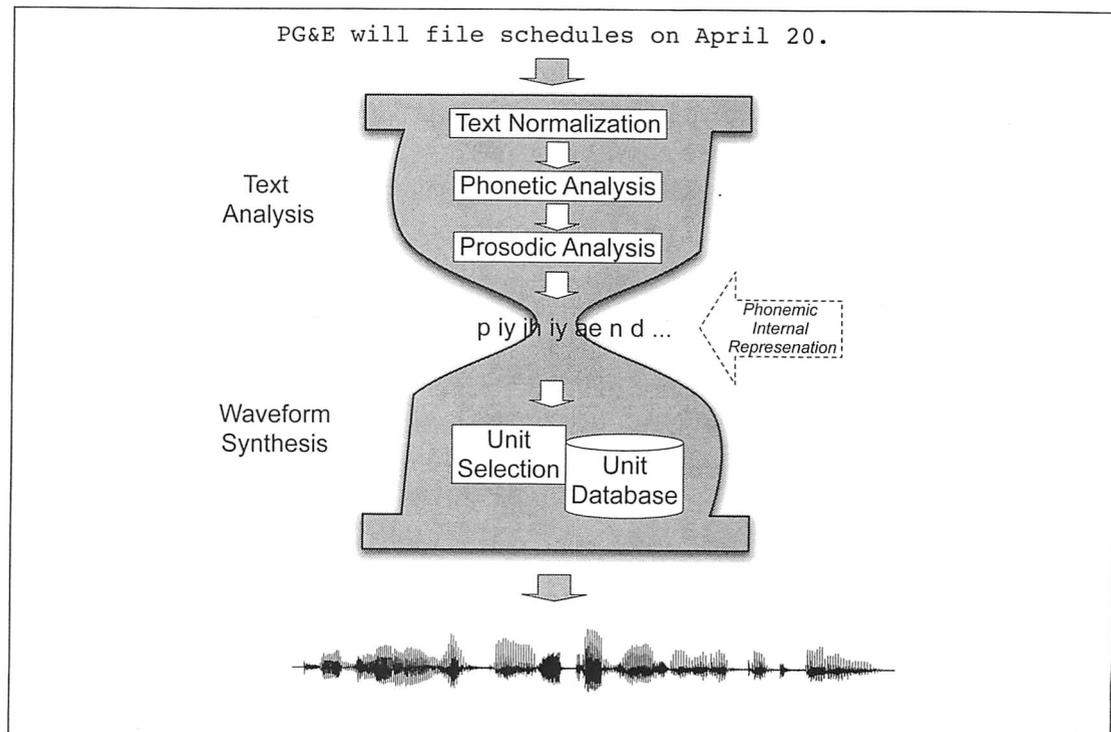


Figure 8.2 The TTS architecture for unit selection (concatenative) speech synthesis.

8.1 Text Normalization

Text normalization

To generate a phonemic internal representation, we must first preprocess or **normalize** the raw text in a variety of ways. We'll need to break the input text into sentences, deal with the idiosyncracies of abbreviations, numbers, and so on. Consider the difficulties in the following text drawn from the Enron corpus (Klimt and Yang, 2004):

He said the increase in credit limits helped B.C. Hydro achieve record net income of about \$1 billion during the year ending March 31. This figure does not include any write-downs that may occur if Powerex determines that any of its customer accounts are not collectible. Cousins, however, was insistent that all debts will be collected: "We continue to pursue monies owing and we expect to be paid for electricity we have sold."

Sentence tokenization

The first task in text normalization is **sentence tokenization**. To segment this paragraph into separate utterances for synthesis, we need to know that the first sentence ends at the period after *March 31*, not at the period of *B.C.* We also need to know that there is a sentence ending at the word *collected*, despite the punctuation being a colon rather than a period. The second normalization task is dealing with **non-standard words**. Non-standard words include numbers, acronyms, abbreviations, and so on. For example, *March 31* needs to be pronounced *March thirty-first*, not *March three one*; *\$1 billion* needs to be pronounced *one billion dollars*, with the word *dollars* appearing after the word *billion*.

8.1.1 Sentence Tokenization

We saw two examples above where sentence tokenization is difficult because sentence boundaries are not always indicated by periods and can sometimes be indicated by punctuation like colons. An additional problem occurs when an abbreviation ends a sentence, in which case the abbreviation-final period is playing a dual role:

- (8.2) He said the increase in credit limits helped B.C. Hydro achieve record net income of about \$1 billion during the year ending March 31.
- (8.3) Cousins, however, was insistent that all debts will be collected: "We continue to pursue monies owing and we expect to be paid for electricity we have sold."
- (8.4) The group included Dr. J. M. Freeman and T. Boone Pickens Jr.

A key part of sentence tokenization is thus period disambiguation; we've seen a simple Perl script for period disambiguation in Chapter 3. Most sentence tokenization algorithms are slightly more complex than this deterministic algorithm, and, in particular, are trained by machine learning methods rather than being hand-built. We do this training by first hand-labeling a training set with sentence boundaries and then applying any supervised machine learning method (decision trees, logistic regression, SVM, etc.) to train a classifier to mark the sentence boundary decisions.

More specifically, we could start by tokenizing the input text into tokens separated by whitespace and then select any token containing one of the three characters `!`, `.` or `?` (or possibly also `:`). After hand-labeling a corpus of such tokens, we train a classifier to make a binary decision (EOS (end-of-sentence) versus not-EOS) on these potential sentence boundary characters inside these tokens.

The success of such a classifier depends on the features that are extracted for the classification. Let's consider some feature templates we might use to disambiguate these **candidate** sentence boundary characters, assuming we have a small amount of training data labeled for sentence boundaries:

- the prefix (the portion of the candidate token preceding the candidate)
- the suffix (the portion of the candidate token following the candidate)
- whether the prefix or suffix is an abbreviation (from a list)
- the word preceding the candidate
- the word following the candidate
- whether the word preceding the candidate is an abbreviation
- whether the word following the candidate is an abbreviation

Consider the following example:

(8.5) ANLP Corp. chairman Dr. Smith resigned.

Given these feature templates, the feature values for the period . in the word Corp. in (8.5) would be

```

PreviousWord = ANLP           NextWord = chairman
Prefix = Corp                Suffix = NULL
PreviousWordAbbreviation = 1 NextWordAbbreviation = 0

```

If our training set is large enough, we can also look for lexical cues about sentence boundaries. For example, certain words may tend to occur sentence-initially or sentence-finally. We can thus add the following features:

- Probability[candidate occurs at end of sentence]
- Probability[word following candidate occurs at beginning of sentence]

Finally, while most of the above features are relatively language independent, we can use language-specific features. For example, in English, sentences usually begin with capital letters, suggesting features like the following:

- case of candidate: Upper, Lower, AllCap, Numbers
- case of word following candidate: Upper, Lower, AllCap, Numbers

Similarly, we can have specific subclasses of abbreviations, such as honorifics or titles (e.g., Dr., Mr., Gen.), corporate designators (e.g., Corp., Inc.), or month names (e.g., Jan., Feb.).

Any machine learning method can be applied to train EOS classifiers. Logistic regression (Section 6.6.2) and decision trees are two common methods; logistic regression may have somewhat higher accuracy, although we have instead shown a decision tree in Fig. 8.3 because it is easier to see how the features are used.

8.1.2 Non-Standard Words

*Non-standard
words*

The second step in text normalization is normalizing **non-standard words** (NSWs). Non-standard words are tokens, like numbers or abbreviations, that need to be expanded into sequences of English words before they can be pronounced.

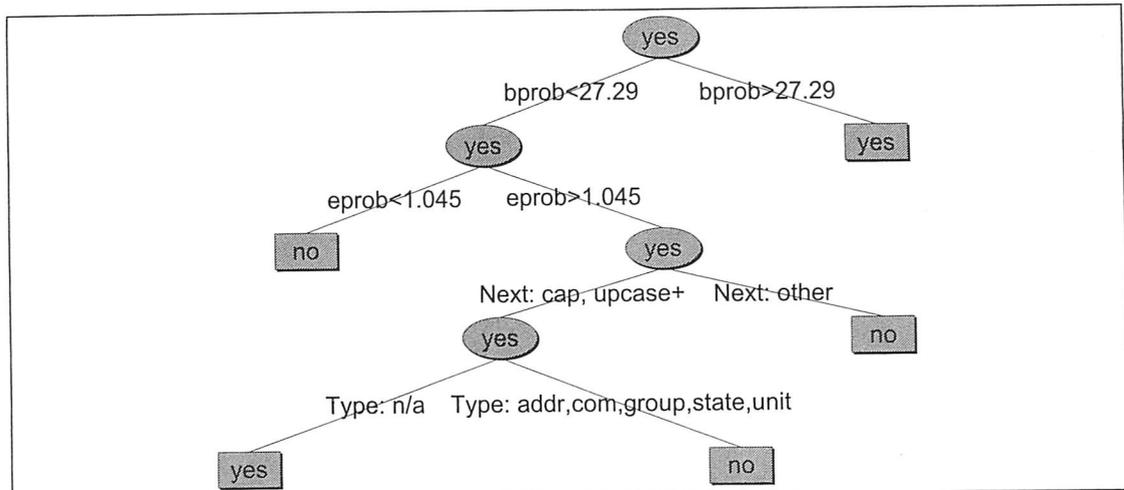


Figure 8.3 A decision tree from Richard Sproat for predicting whether a period “.” is an end of sentence (YES) or not an end-of-sentence (NO). We use features like the log likelihood of the current word being the beginning of a sentence (*bprob*), the previous word being an end-of-sentence (*eprob*), the capitalization of the next word, and the abbreviation subclass (company, state, unit of measurement).

Non-standard words are difficult because they are often ambiguous. For example, the number *1750* can be spoken in at least four different ways, depending on the context:

seventeen fifty: (in “*The European economy in 1750*”)
 one seven five zero: (in “*The password is 1750*”)
 seventeen hundred and fifty: (in “*1750 dollars*”)
 one thousand, seven hundred, and fifty: (in “*1750 dollars*”)

Similar ambiguities occur for Roman numerals like *IV* (which can be pronounced *four*, *fourth*, or as the letters *I V* (meaning “intravenous”)), or *2/3*, which can be *two thirds* or *February third* or *March second* or *two slash three*.

Some non-standard words are composed of letters, such as **abbreviations**, **letter sequences**, and **acronyms**. Abbreviations are generally pronounced as if **expanded**; thus, *Wed* is pronounced *Wednesday*, and *Jan 1* is pronounced *January first*. **Letter sequences** like *UN*, *DVD*, *PC*, and *IBM* are pronounced letter by letter in a sequence; *IBM* is thus pronounced [ay b iy eh m]. **Acronyms** like *IKEA*, *MoMA*, *NASA*, and *UNICEF* are pronounced as if they were words; *MoMA* is pronounced [m ow m ax]. Ambiguity occurs here as well; should *Jan* be read as a word (the name *Jan*) or expanded as the month *January*? Different types of numeric and alphabetic non-standard words are summarized in Fig. 8.4.

Each of the types has a particular realization (or realizations). For example, a year (NYER) is generally read in the **paired** method, in which each pair of digits is pronounced as an integer (e.g., *seventeen fifty* for *1750*), while a U.S. zip code (NZIP) is generally read in the **serial** method, as a sequence of single digits (e.g., *nine four one one zero* for *94110*). The type *BMONEY* deals with the idiosyncracies of expressions like *\$3.2 billion*, which must be read out with the word *dollars* at the end, as *three point two billion dollars*. For the alphabetic NSWs, we have the class *EXPN* for abbreviations like *N.Y.* that are expanded,

Paired digits

Serial digits

ALPHA	EXPN	abbreviation	<i>adv, N.Y., mph, gov't</i>
	LSEQ	letter sequence	<i>DVD, D.C., PC, UN, IBM,</i>
	ASWD	read as word	<i>IKEA, unknown words/names</i>
NUMBERS	NUM	number (cardinal)	<i>12, 45, 1/2, 0.6</i>
	NORD	number (ordinal)	<i>May 7, 3rd, Bill Gates III</i>
	NTEL	telephone (or part of)	<i>212-555-4523</i>
	NDIG	number as digits	<i>Room 101</i>
	NIDE	identifier	<i>747, 386, 15, pc110, 3A</i>
	NADDR	number as street address	<i>747, 386, 15, pc110, 3A</i>
	NZIP	zip code or PO Box	<i>91020</i>
	NTIME	a (compound) time	<i>3.20, 11:45</i>
	NDATE	a (compound) date	<i>2/28/05, 28/02/05</i>
	NYER	year(s)	<i>1998, 80s, 1900s, 2008</i>
	MONEY	money (US or other)	<i>\$3.45, HK\$300, Y20,200, \$200K</i>
	BMONEY	money tr/m/billions	<i>\$3.45 billion</i>
	PRCT	percentage	<i>75% 3.4%</i>

Figure 8.4 Some types of non-standard words in text normalization, selected from Table 1 of Sproat et al. (2001); not listed are types for URLs, emails, and some complex uses of punctuation.

LSEQ for acronyms pronounced as letter sequences, and ASWD for acronyms pronounced as if they were words.

Dealing with non-standard words requires at least three steps: **tokenization** to separate and identify potential non-standard words, **classification** to label them with a type from Fig. 8.4, and **expansion** to convert each type into a string of standard words.

In the tokenization step, we can tokenize the input by whitespace and then assume that any word not in the pronunciation dictionary is a non-standard word. More sophisticated tokenization algorithms would also deal with the fact that some dictionaries already contain some abbreviations. The CMU dictionary, for example, contains abbreviated (and hence incorrect) pronunciations for *st*, *mr*, *mrs*, as well as day and month abbreviations like *mon*, *tues*, *nov*, *dec*, etc. Thus, in addition to unseen words, we also need to label any of these acronyms and also single-character tokens as potential non-standard words. Tokenization algorithms also need to split words that are combinations of two tokens, like *2-car* or *RVing*. Words can be split by simple heuristics, such as splitting at dashes, or at changes from lower case to upper case.

The next step is assigning an NSW type; many types can be detected with simple regular expressions. For example, NYER could be detected by the following regular expression:

$$/(1[89][0-9][0-9])|(20[0-9][0-9])/$$

Other classes might be harder to write rules for, so a more powerful option is to use a machine learning classifier with many features.

To distinguish between the alphabetic ASWD, LSEQ and EXPN classes, for example we might want features over the component letters. Thus short, all-capital words (*IBM*, *US*) might be LSEQ, longer all-lowercase words with a single-quote (*gov't*, *cap'n*) might be EXPN, and all-capital words with multiple vowels (*NASA*, *IKEA*) might be more likely to be ASWD.

Another useful feature is the identity of neighboring words. Consider ambiguous strings like *3/4*, which can be a NUM (*three-fourths*) or an NDATE (*march third*). NDATE might be preceded by the word *on*, followed by the word *of*, or have the word *Monday* somewhere in the surrounding words. By contrast, NUM examples might be preceded by another number or followed by words like *mile* and *inch*. Similarly, Roman numerals like *VII* tend to be NORD (*seven*) when preceded by *Chapter*, *part*, or *Act*, but NUM (*seventh*) when the words *king* or *Pope* occur in the neighborhood. These context words can be chosen as features by hand or can be learned by machine learning techniques like the **decision list** algorithm of Chapter 20.

We can achieve the most power by building a single machine learning classifier that combines all the above ideas. For example, the NSW classifier of Sproat et al. (2001) uses 136 features, including letter-based features like '*all-upper-case*', '*has-two-vowels*', '*contains-slash*', and '*token-length*', as well as binary features for the presence of certain words like *Chapter*, *on*, or *king* in the surrounding context. Sproat et al. (2001) also included a rough-draft, rule-based classifier, which used hand-written regular expression to classify many of the number NSWs. The output of this rough-draft classifier was used as just another feature in the main classifier.

To build such a main classifier, we need a hand-labeled training set in which each token has been labeled with its NSW category; one such hand-labeled database was produced by Sproat et al. (2001). Given the labeled training set, we can apply any supervised machine learning algorithm, such as the logistic regression or decision tree algorithms discussed earlier. For each observed token o_i we extract the features discussed above. Then we train the classifier to use these features to try to predict the hand-labeled NSW category from Fig. 8.4.

The third step in dealing with NSWs is expansion into ordinary words. One NSW type, EXPN, is quite difficult to expand. The type includes abbreviations and acronyms like *NY*. Generally these must be expanded with the help of an abbreviation dictionary, and any ambiguities can be dealt with by the homonym disambiguation algorithms discussed in the next section.

Expansion of the other NSW types is generally deterministic. Many expansions are trivial; for example, LSEQ expands to a sequence of words, one for each letter, ASWD expands to itself, NUM expands to a sequence of words representing the cardinal number, NORD expands to a sequence of words representing the ordinal number, and NDIG and NZIP both expand to a sequence of words, one for each digit.

Other types are slightly more complex; NYER expands to two pairs of digits, unless the year ends in *00*, in which case the four years are pronounced as a cardinal number (*2000* as *two thousand*) or in the **hundreds** method (e.g., *1800* as *eighteen hundred*). NTEL can be expanded just as a sequence of digits; alternatively, the last four digits can be read as **paired digits**, in which each pair is read as an integer. It is also possible to read them in a form known as **trailing unit**, in which the digits are read serially until the last non-zero digit, which is pronounced followed by the appropriate unit (e.g., *876-5000* as *eight seven six five thousand*). The expansions of NDATE, MONEY, and NTIME are left as Exercises 8.1–8.4 for the reader.

Of course, many of these expansions are dialect-specific. In Australian English, the sequence *33* in a telephone number is generally read *double three*. Other languages also present additional difficulties in non-standard word normalization. In

Hundreds digits

Trailing unit digits

French or German, for example, in addition to the above issues, normalization may depend on morphological properties. In French, the phrase *l fille* ('one girl') is normalized to *une fille*, but *l garçon* ('one boy') is normalized to *un garçon*. Similarly, in German, *Heinrich IV* ('Henry IV') can be normalized to *Heinrich der Vierte*, *Heinrich des Vierten*, *Heinrich dem Vierten*, or *Heinrich den Vierten* depending on the grammatical case of the noun (Demberg, 2006).

8.1.3 Homograph Disambiguation

The goal of our NSW algorithms in the previous section was to determine which sequence of standard words to pronounce for each NSW. But sometimes determining how to pronounce even standard words is difficult. This is particularly true for **homographs**, which are words with the same spelling but different pronunciations. Here are some examples of the English homographs *use*, *live*, and *bass*:

Homograph

- (8.6) It's no use (/y uw s/) to ask to use (/y uw z/) the telephone.
 (8.7) Do you live (/l ih v/) near a zoo with live (/l ay v/) animals?
 (8.8) I prefer bass (/b ae s/) fishing to playing the bass (/b ey s/) guitar.

French homographs include *fil*s (which has two pronunciations [fis] 'son' versus [fil] 'thread'), or the multiple pronunciations for *fier* ('proud' or 'to trust'), and *est* ('is' or 'East') (Divay and Vitale, 1997).

Luckily for the task of homograph disambiguation, the two forms of homographs in English (as well as in similar languages like French and German) tend to have different parts of speech. For example, the two forms of *use* above are (respectively) a noun and a verb, while the two forms of *live* are (respectively) a verb and a noun. Figure 8.5 shows some interesting systematic relations between the pronunciation of some noun-verb and adj-verb homographs. Indeed, Liberman and Church (1992) showed that many of the most frequent homographs in 44 million words of AP newswire are disambiguatable just by using part of speech (the most frequent 15 homographs in order are *use*, *increase*, *close*, *record*, *house*, *contract*, *lead*, *live*, *lives*, *protest*, *survey*, *project*, *separate*, *present*, *read*).

Final voicing		Stress shift		-ate final vowel	
N (/s/)	V (/z/)	N (init. stress)	V (fin. stress)	N/A (final /ax/)	V (final /ey/)
use	y uw s	record	r eh1 k axr0 d	estimate	eh s t ih m ax t
close	k l ow s	insult	ih1 n s ax0 l t	separate	s eh p ax r ax t
house	h aw s	object	aa1 b j eh0 k t	moderate	m aa d ax r ax t
	y uw z		r ix0 k ao1 r d		eh s t ih m ey t
	k l ow z		ix0 n s ah1 l t		s eh p ax r ey t
	h aw z		ax0 b j eh1 k t		m aa d ax r ey t

Figure 8.5 Some systematic relationships between homographs: final consonant (noun /s/ versus verb /z/), stress shift (noun initial stress versus verb final stress), and final vowel weakening in *-ate* noun/adjs.

Because knowledge of part of speech is sufficient to disambiguate many homographs, in practice we perform homograph disambiguation by storing distinct pronunciations for these homographs labeled by part of speech and then running a *part-of-speech* tagger to choose the pronunciation for a given homograph in context.

There are a number of homographs, however, for which both pronunciations have the same part of speech. We saw two pronunciations for *bass* (fish versus instrument)

above. Other examples of these include *lead* (because there are two noun pronunciations, /l iy d/ (a leash or restraint) and /l eh d/ (a metal)). We can also think of the task of disambiguating certain abbreviations (mentioned earlier as NSW disambiguation) as homograph disambiguation. For example, *Dr.* is ambiguous between *doctor* and *drive*, and *St.* between *Saint* or *street*. Finally, there are some words that differ in capitalizations like *polish/Polish*, which are homographs only in situations like sentence beginnings or all-capitalized text.

In practice, these latter classes of homographs that cannot be resolved by part of speech are often ignored in TTS systems. Alternatively, we can attempt to resolve them by using the word sense disambiguation algorithms, like the **decision-list** algorithm of Yarowsky (1997), which we will introduce in Chapter 20.

8.2 Phonetic Analysis

The next stage in synthesis is to produce a pronunciation for each word in the normalized word strings from text analysis. The most important component here is a large pronunciation dictionary. Dictionaries alone turn out to be insufficient, however, since running text always contains words that don't appear in the dictionary. For example Black et al. (1998) used a British English dictionary, the OALD lexicon on the first section of the Penn Wall Street Journal Treebank. Of the 39,923 words (tokens) in this section, 1775 word tokens (4.6%) were not in the dictionary, of which 943 are unique (i.e., 943 types). The distributions of these unseen word tokens was as follows:

Names	Unknown	Typos and other
1360	351	64
76.6%	19.8%	3.6%

Thus, the two main areas in which dictionaries need to be augmented are in dealing with names and with other unknown words. The next three sections discuss each of these in order: dictionaries, names, and grapheme-to-phoneme rules for other unknown words.

8.2.1 Dictionary Lookup

Phonetic dictionaries were introduced in Section 7.5. One of the most widely used for TTS is the freely available CMU Pronouncing Dictionary (CMU, 1993), which has pronunciations for about 120,000 words. The pronunciations are roughly phonemic, from a 39-phone ARPAbet-derived phoneme set. Phonemic transcriptions means that instead of marking surface reductions like the reduced vowels [ax] or [ix], CMUdict marks each vowel with a stress tag: 0 (unstressed), 1 (primary stress), or 2 (secondary stress). Thus, (non-diphthong) vowels with 0 stress generally correspond to [ax] or [ix]. Most of the words have only a single pronunciation, but about 8,000 of them have two or even three pronunciations and some kinds of phonetic reductions are marked in these pronunciations. The dictionary is not syllabified, although the nucleus is implicitly marked by the (numbered) vowel. Figure 8.6 shows some sample pronunciations.

<i>ANTECEDENTS</i>	AE2 N T IH0 S IY1 D AH0 N T S	<i>PAKISTANI</i>	P AE2 K IH0 S T AE1 N IY0
<i>CHANG</i>	CH AE1 NG	<i>TABLE</i>	T EY1 B AH0 L
<i>DICTIONARY</i>	D IH1 K SH AH0 N EH2 R IY0	<i>TROTSKY</i>	T R AA1 T S K IY2
<i>DINNER</i>	D IH1 N ER0	<i>WALTER</i>	W AO1 L T ER0
<i>LUNCH</i>	L AH1 N CH	<i>WALTZING</i>	W AO1 L T S IH0 NG
<i>MCFARLAND</i>	M AH0 K F AA1 R L AH0 N D	<i>WALTZING(2)</i>	W AO1 L S IH0 NG

Figure 8.6 Some sample pronunciations from the CMU Pronouncing Dictionary.

The CMU dictionary was designed for speech recognition rather than synthesis uses; thus, it does not specify which of the multiple pronunciations to use for synthesis, does not mark syllable boundaries, and because it capitalizes the dictionary head-words, does not distinguish between, for example, *US* and *us* (the form *US* has the two pronunciations [AH1 S] and [Y UW1 EH1 S]).

The 110,000 word UNISYN dictionary, freely available for research purposes, resolves many of these issues as it was designed specifically for synthesis (Fitt, 2002). UNISYN gives syllabifications, stress, and some morphological boundaries. Furthermore, pronunciations in UNISYN can also be read off in any of dozens of dialects of English, including General American, RP British, Australia, and so on. The UNISYN uses a slightly different phone set; here are some examples:

```
going:      { g * ou }.> i ng >
antecedents: { * a n . t ^ i . s ~ ii . d n ! t } > s >
dictionary:  { d * i k . sh @ . n ~ e . r ii }
```

8.2.2 Names

The distribution of unknown words discussed on page 257 above indicated the importance of names, including for example personal names (first names and surnames), geographical names (city, street, and other place names), and commercial names (company and product names). Considering only personal names, Spiegel (2003) estimates that there are about two million different surnames and 100,000 first names just for the United States. Two million is a very large number: an order of magnitude more than the entire size of the CMU dictionary. For this reason, most large-scale TTS systems include a large name-pronunciation dictionary. As we saw in Fig. 8.6 the CMU dictionary itself contains a wide variety of names; in particular, it includes the pronunciations of the most frequent 50,000 surnames from an old Bell Lab estimate of US personal name frequency, as well as 6,000 first names.

How many names are sufficient? Liberman and Church (1992) found that a dictionary of 50,000 names covered 70% of the name tokens in 44 million words of AP newswire. Interestingly, many of the remaining names (up to 97.43% of the tokens in their corpus) could be accounted for by simple modifications of these 50,000 names, such as adding stress-neutral suffixes to names like *Walter* or *Lucas* to produce *Walters* or *Lucasville*. Other pronunciations might be created by rhyme analogy. If we have the pronunciation for the name *Trotsky*, but not the name *Plotsky*, we can replace the initial /tr/ from *Trotsky* with initial /pl/ to derive a pronunciation for *Plotsky*.

Techniques such as this, including morphological decomposition, analogical formation, and mapping of unseen names to spelling variants already in the dictionary

(Fackrell and Skut, 2004), have achieved some success in name pronunciation. In general, however, name pronunciation is still difficult. Many modern systems deal with unknown names via the grapheme-to-phoneme methods described in the next section, often by building two predictive systems, one for names and one for non-names. Spiegel (2003, 2002) summarizes many more issues in proper name pronunciation.

8.2.3 Grapheme-to-Phoneme Conversion

Once we have expanded non-standard words and looked them all up in a pronunciation dictionary, we need to pronounce the remaining, unknown words. The process of converting a sequence of letters into a sequence of phones is called **grapheme-to-phoneme** conversion, sometimes shortened **g2p**. The job of a grapheme-to-phoneme algorithm is thus to convert a letter string like *cake* into a phone string like [K EY K].

The earliest such algorithms were rules written by hand in the Chomsky-Halle rewrite rule format of Eq. 7.1 in Chapter 7. These are often called **letter-to-sound** or LTS rules, and they are sometimes still used. LTS rules are applied in order, with later (default) rules applying only if the context for earlier rules is not applicable. A simple pair of rules for pronouncing the letter *c* might be as follows:

$$c \rightarrow [k] / _ \{a,o\}V \quad ; \text{context dependent} \quad (8.9)$$

$$c \rightarrow [s] \quad ; \text{context independent} \quad (8.10)$$

Actual rules must be much more complicated (for example *c* can also be pronounced [ch] in *cello* or *concerto*). Even more complex are rules for assigning stress, which are famously difficult for English. Consider just one of the many stress rules from Allen et al. (1987), whereby the symbol *X* represents all possible syllable onsets:

$$V \rightarrow [+stress] / X _ C^* \{V_{\text{short}} C C? | V\} \{V_{\text{short}} C^* | V\} \quad (8.11)$$

This rule represents the following two situations:

1. Assign 1-stress to the vowel in a syllable followed by a weak syllable followed by a morpheme-final syllable with a short vowel and 0 or more consonants (e.g., *difficult*).
2. Assign 1-stress to the vowel in a syllable preceding a weak syllable followed by a vowel that is morpheme-final (e.g., *oregano*).

While some modern systems still use such complex hand-written rules, most systems achieve higher accuracy by relying instead on automatic or semi-automatic methods based on machine learning. This modern probabilistic grapheme-to-phoneme problem was first formalized by Lucassen and Mercer (1984). Given a letter sequence *L*, we are searching for the most probable phone sequence *P*:

$$\hat{P} = \underset{P}{\operatorname{argmax}} P(P|L) \quad (8.12)$$

The probabilistic method assumes a training set and a test set; both sets are lists of words from a dictionary, with a spelling and a pronunciation for each word. The next subsections show how a classifier can be trained to estimate this probability $P(P|L)$ and applied to produce the pronunciation for an unseen word.

Grapheme-to-
phoneme

Letter-to-sound

Finding a Letter-to-Phone Alignment for the Training Set

Most letter-to-phone algorithms assume that we have an **alignment** that tells us which phones align with each letter. We'll need this alignment for each word in the training set. A letter can align to multiple phones (e.g., *x* often aligns to *k s*), or to no phones at all, like the final letter of *cake* in the following alignment:

L:	c	a	k	e
P:	K	EY	K	ε

One method for finding such a letter-to-phone alignment is the semi-automatic method of Black et al. (1998). Their algorithm is semi-automatic because it relies on a hand-written list of the **allowable** phones that can realize each letter. Here are allowable lists for the letters *c* and *e*:

c: k ch s sh t-s ε

e: ih iy er ax ah eh ey uw ay ow y-uw oy aa ε

To produce an alignment for each word in the training set, we take this allowables list for all the letters, and for each word in the training set, we find all alignments between the pronunciation and the spelling that conform to the allowables list. From this large list of alignments, we compute, by summing over all alignments for all words, the total count for each letter being aligned to each phone (or multi-phone or ε). From these counts we can normalize to get for each phone p_i and letter l_j a probability $P(p_i|l_j)$:

$$P(p_i|l_j) = \frac{\text{count}(p_i, l_j)}{\text{count}(l_j)} \quad (8.13)$$

We can now take these probabilities and realign the letters to the phones, using the Viterbi algorithm to produce the best (Viterbi) alignment for each word, where the probability of each alignment is just the product of all the individual phone/letter alignments. The result is a single good alignment A for each training pair (P, L) .

Choosing the Best Phone String for the Test Set

Given a new word w , we now need to map its letters into a phone string. We'll train a machine learning classifier on the aligned training set. The classifier will look at a letter of the word and generate the most probable phone. Obviously, we can do a better job of predicting the phone if we look at a window of surrounding letters; for example, consider the letter *a*. In the word *cat*, the *a* is pronounced *AE*. But in our word *cake*, *a* is pronounced *EY* because *cake* has a final *e*; thus, knowing whether there is a final *e* is a useful feature. Typically, we look at the k previous letters and the k following letters.

Another useful feature would be the correct identity of the previous phone. Knowing this would allow us to get some phonotactic information into our probability model. Of course, we can't know the true identity of the previous phone, but we can approximate this by looking at the previous phone that was predicted by our model. To do this, we can run our classifier left to right, generating phones one by one.

In summary, in the most common classifier model, the probability of each phone p_i is estimated from a window of k previous and k following letters, as well as the most recent k phones that were previously produced.

Figure 8.7 shows a sketch of this left-to-right process, indicating the features that the classifier would use to choose a phone for the letter *s* in the word *Jurafsky*. We can integrate stress prediction into phone prediction by augmenting our set of phones with stress information, for example having two copies of each vowel (e.g., AE and AE1) or even the three levels of stress AE0, AE1, and AE2 used in the CMU lexicon. Other useful features include the part-of-speech tag of the word (most part-of-speech taggers provide an estimate of the part-of-speech tag even for unknown words) and facts such as whether the previous vowel was stressed or even using classes of letters (corresponding roughly to consonants, vowels, liquids, and so on).

Liaison

For some languages, we also need to know features about the following word. For example, French has a phenomenon called **liaison**, in which the realization of the final phone of some words depends on whether there is a next word and whether that word starts with a consonant or a vowel. For example, the French word *six* can be pronounced [sis] (in *j'en veux six* 'I want six'), [siz] (*six enfants* 'six children'), [si] (*six filles* 'six girls').

Finally, most synthesis systems build two separate grapheme-to-phoneme classifiers, one for unknown personal names, and one for other unknown words. For pronouncing personal names it turns out to be helpful to use additional features that indicate which foreign language the names originally came from. Such features could be the output of a foreign-language classifier based on letter sequences (different languages have characteristic letter *N*-gram sequences).

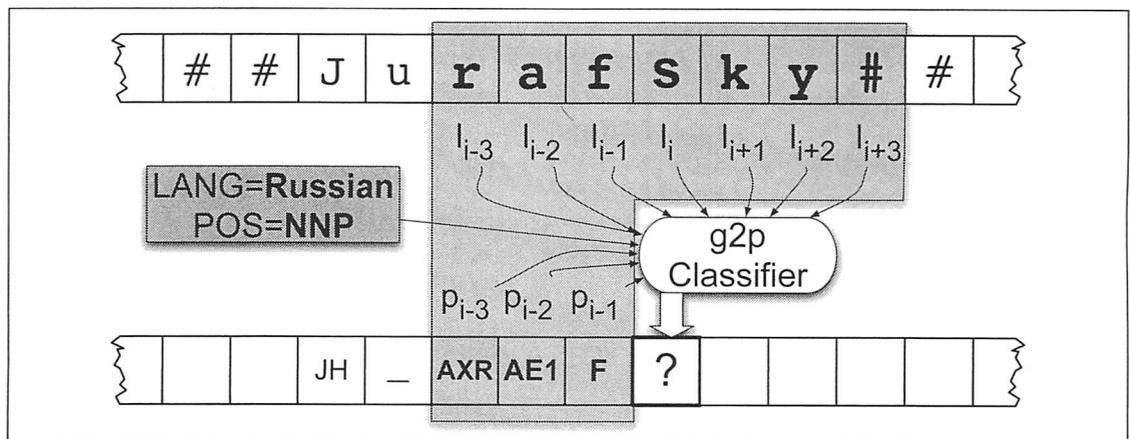


Figure 8.7 The left-to-right process of converting graphemes to phonemes, showing the decision for the letter *s*. The features are shown in shading, using a context window $k = 3$; real TTS systems use window sizes of 5 or larger.

Decision trees and logistic regression are conditional classifiers, computing the phoneme string that has the highest conditional probability given the grapheme sequence. More recent grapheme-to-phoneme conversion makes use of a joint classifier in which the hidden state is a combination of phone and grapheme called a **graphone**; see the Historical Notes section at the end of the chapter for references.

8.3 Prosodic Analysis

Prosody

The final stage of linguistic analysis is prosodic analysis. In poetry, the word **prosody** refers to the study of the metrical structure of verse. In linguistics and language processing, however, we use the term **prosody** to mean the study of the intonational and rhythmic aspects of language. More technically, prosody has been defined by Ladd (1996) as the “use of suprasegmental features to convey sentence-level pragmatic meanings”. The term **suprasegmental** means above and beyond the level of the segment or phone. The term refers especially to the uses of acoustic features like **F0**, **duration**, and **energy** independently of the phone string.

Suprasegmental

By **sentence-level pragmatic meaning**, Ladd is referring to a number of kinds of meaning that have to do with the relation between a sentence and its discourse or external context. For example, prosody can be used to mark **discourse structure or function**, like the difference between statements and questions, or the way that a conversation is structured into segments or subdialogs. Prosody is also used to mark **saliency**, such as indicating that a particular word or phrase is important or salient. Finally, prosody is heavily used for affective and emotional meaning, such as expressing happiness, surprise, or anger.

In the next sections we introduce the three aspects of prosody, each of which is important for speech synthesis: **prosodic structure**, **prosodic prominence**, and **tune**. Prosodic analysis generally proceeds in two parts. First, we compute an abstract representation of the prosodic structure, prominence, and tune of the text. For unit selection synthesis, this is all we need to do in the text analysis component. For diphone and HMM synthesis, we have one further step, which is to predict **duration** and **F0** values from these prosodic structures.

8.3.1 Prosodic Structure

Prosodic phrasing

Spoken sentences have prosodic structure in the sense that some words seem to group naturally together and some words seem to have a noticeable break or disjuncture between them. Prosodic structure is often described in terms of **prosodic phrasing**, meaning that an utterance has a prosodic phrase structure in a similar way to it having a syntactic phrase structure. For example, in the sentence *I wanted to go to London, but could only get tickets for France* there seem to be two main **intonation phrases**, their boundary occurring at the comma. Furthermore, in the first phrase, there seems to be another set of lesser prosodic phrase boundaries (often called **intermediate phrases**) that split up the words as *I wanted | to go | to London*.

Intonation phrase

Intermediate phrase

Prosodic phrasing has many implications for speech synthesis; the final vowel of a phrase is longer than usual, we often insert a pause after an intonation phrases, and, as we discuss in Section 8.3.6, there is often a slight drop in F0 from the beginning of an intonation phrase to its end, and F0 then resets at the beginning of a new intonation phrase.

Practical phrase boundary prediction is generally treated as a binary classification task, where we are given a word and we have to decide whether to put a prosodic boundary after it. A simple model for boundary prediction can be based on determinis-

tic rules. A very high precision rule is the one we saw for sentence segmentation: insert a boundary after punctuation. Another commonly used rule inserts a phrase boundary before a function word following a content word.

More sophisticated models are based on machine learning classifiers. To create a training set for classifiers, we first choose a corpus and then mark every prosodic boundary in the corpus. One way to do this prosodic boundary labeling is to use an intonational model like ToBI or Tilt (see Section 8.3.4) and have human labelers listen to speech and label the transcript with the boundary events defined by the theory. Because prosodic labeling is extremely time-consuming, however, a text-only alternative is often used. In this method, a human labeler looks only at the text of the training corpus, ignoring the speech. The labeler marks any juncture between words at which a prosodic boundary might legitimately occur if the utterance were spoken.

Given a labeled training corpus, we can train a classifier to make a binary (boundary vs. no boundary) decision at every juncture between words (Wang and Hirschberg, 1992; Ostendorf and Veilleux, 1994; Taylor and Black, 1998).

Features that are commonly used in classification include the following:

- **Length features:** Phrases tend to be of roughly equal length, so we can use various features that hint at phrase length (Bachenko and Fitzpatrick, 1990; Grosjean et al., 1979; Gee and Grosjean, 1983).
 - The total number of words and syllables in utterance
 - The distance of the juncture from the beginning and end of the sentence (in words or syllables)
 - The distance in words from the last punctuation mark
- **Neighboring part of speech and punctuation:**
 - The part-of-speech tags for a window of words around the juncture. Generally, the two words before and after the juncture are used.
 - The type of following punctuation

There is also a correlation between prosodic structure and **syntactic structure** (Price et al., 1991). Thus, robust parsers like Collins (1997) can be used to label the sentence with rough syntactic information, from which we can extract syntactic features such as the size of the biggest syntactic phrase that ends with this word (Ostendorf and Veilleux, 1994; Koehn et al., 2000). We will introduce syntactic structure and parsing in Chapters 12–14.

8.3.2 Prosodic Prominence

In any spoken utterance, some words sound more **prominent** than others. Prominent words are perceptually more salient to the listener; speakers make a word more salient in English by saying it louder, saying it slower (so it has a longer duration), or by varying F0 during the word, making it higher or more variable.

We generally capture the core notion of prominence by associating a linguistic marker with prominent words, a marker called **pitch accent**. Words that are prominent are said to **bear** (be associated with) a pitch accent. Pitch accent is thus part of the phonological description of a word in context in a spoken utterance.

Prominence

Pitch accent

Pitch accent is related to **stress**, which we discussed in Chapter 7. The stressed syllable of a word is where pitch accent is realized. In other words, if a speaker decides to highlight a word by giving it a pitch accent, the accent will appear on the stressed syllable of the word. The following example shows accented words in capital letters, with the stressed syllable bearing the accent (the louder, longer syllable) in boldface:

(8.14) I'm a little **SURPRISED** to hear it **CHARACTERIZED** as **UPBEAT**.

Nuclear accent

We generally need more fine-grained distinctions than just a binary distinction between accented and unaccented words. For example, the last accent in a phrase generally is perceived as being more prominent than the other accents. This prominent last accent is called the **nuclear accent**. Emphatic accents are generally used for semantic purposes, for example, to indicate that a word is the **semantic focus** of the sentence (see Chapter 21) or that a word is contrastive or otherwise important in some way. Such emphatic words are often written IN CAPITAL LETTERS or with ****stars**** around them in SMS or email or *Alice in Wonderland*; here's an example from the latter:

(8.15) "I know **SOMETHING** interesting is sure to happen," she said to herself.

Some words can also be **less prominent** than usual, such as function words, which are often phonetically very **reduced** (see Chapter 7). Accents can also differ according to the **tune** associated with them. Accents with particularly high pitch, for example, have different functions than those with particularly low pitch; we show how this is modeled in the ToBI model in Section 8.3.4.

Ignoring tune for the moment, we can summarize by saying that speech synthesis systems can use as many as four levels of prominence: **emphatic accent**, **pitch accent**, **unaccented**, and **reduced**. In practice, however, many implemented systems make do with a subset of only two or three of these levels.

With two levels, pitch accent prediction is a binary classification task, for which we must decide whether a given word is accented or not. Since in general, informative words (content words, and especially those that are new or unexpected) tend to bear accent (Ladd, 1996; Bolinger, 1972), the simplest accent prediction system is to just accent all content words and no function words.

Better models might seem to require sophisticated semantic knowledge, for example, to understand if a word is new or old in the discourse, whether it is being used contrastively, and exactly how much new information a word contains, and such information was indeed used in early systems (Hirschberg, 1993). But Hirschberg and others showed better prediction by using simple, robust features that correlate with these sophisticated semantics.

tf-idf

For example, the fact that new or unpredictable information tends to be accented can be modeled with robust features like *N*-grams or tf-idf (Pan and Hirschberg, 2000; Pan and McKeown, 1999). The unigram probability of a word $P(w_i)$ and its bigram probability $P(w_i|w_{i-1})$ both correlate with accent; the more probable a word, the less likely it is to be accented. Similarly, an information-retrieval measure known as **tf-idf** (Term-Frequency/Inverse-Document Frequency; see Chapter 23) is a useful accent predictor. Tf-idf captures the semantic importance of a word in a particular document d , by downgrading words that appear in many different documents in a large *N*-document background corpus. One of the many versions of tf-idf can be expressed formally as

follows, assuming $tf_{i,j}$ is the frequency of w_i in the document d_j and n_i is the total number of documents in the corpus that contain w_i :

$$\begin{aligned} \text{idf}_i &= \log \left(\frac{N}{n_i} \right) \\ \text{tf-idf}_{i,j} &= \text{tf}_{i,j} \times \text{idf}_i \end{aligned} \quad (8.16)$$

Accent ratio For words that have been seen enough times in a training set, we can use the **accent ratio** feature, which models a word's individual probability of being accented. The accent ratio of a word is just the probability of the word being accented (if this probability is significantly different from 0.5, and is 0.5 otherwise). More formally,

$$\text{AccentRatio}(w) = \begin{cases} \frac{k}{N} & \text{if } B(k, N, 0.5) \leq 0.05 \\ 0.5 & \text{otherwise} \end{cases}$$

where N is the total number of times the word w occurred in the training set, k is the number of times it was accented, and $B(k, n, 0.5)$ is the probability (under a binomial distribution) that there are k successes in n trials if the probability of success and failure is equal (Nenkova et al., 2007; Yuan et al., 2005).

Features like part of speech, N -grams, tf-idf, and accent ratio can then be combined in a classifier to predict accents. While these robust features work relatively well, a number of problems in accent prediction still remain the subject of research.

For example, it is difficult to predict which of the two words should be accented in adjective-noun or noun-noun compounds. Some regularities do exist; for example, adjective-noun combinations like *new truck* are likely to have accent on the right word (*new TRUCK*), while noun-noun compounds like *TREE surgeon* are likely to have accent on the left. But exceptions to these rules make accent prediction in noun compounds quite complex. For example, the noun-noun compound *APPLE cake* has the accent on the first word, while the noun-noun compound *apple PIE* or *city HALL* both have the accent on the second word (Lieberman and Sproat, 1992; Sproat, 1994, 1998a).

Clash
Lapse Another complication has to do with rhythm; in general, speakers avoid putting accents too close together (a phenomenon known as **clash**) or too far apart (**lapse**). Thus, *city HALL* and *PARKING lot* combine as *CITY hall PARKING lot* with the accent on *HALL* shifting forward to *CITY* to avoid the clash with the accent on *PARKING* (Lieberman and Prince, 1977),

Some of these rhythmic constraints can be captured by using sequence modeling methods such as the sequence classifier shown in Fig. 8.7, in which we run a classifier left to right through a sentence, using the output of the previous word as a feature. We can also use more sophisticated machine learning models like MEMMs (Chapter 6) or conditional random fields (CRFs) (Gregory and Altun, 2004).

8.3.3 Tune

Tune Two utterances with the same prominence and phrasing patterns can still differ prosodically by having different **tunes**. The **tune** of an utterance is the rise and fall of its F0 over time. A very obvious example of tune is the difference between statements

and yes-no questions in English. The same sentence can be said with a final rise in F0 to indicate a yes-no question, or a final fall in F0 to indicate a declarative intonation. Figure 8.8 shows the F0 track of the same words spoken as a question or a statement. Note that the question rises at the end; this is often called a **question rise**. The falling intonation of the statement is called a **final fall**.

Question rise
Final fall

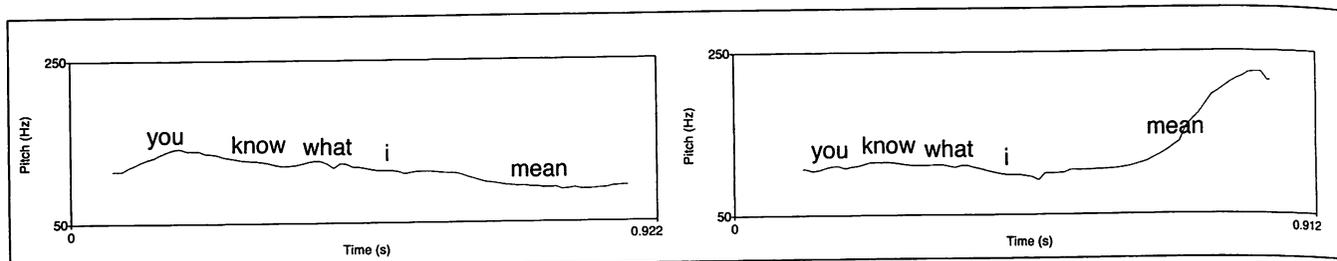


Figure 8.8 The same text read as the statement *You know what I mean* (on the left) and as a question *You know what I mean?* (on the right). Notice that yes-no question intonation in English has a sharp final rise in F0.

It turns out that English makes wide use of tune to express meaning. Besides this well-known rise for yes-no questions, an English phrase containing a list of nouns separated by commas often has a short rise called a **continuation rise** after each noun. Other examples include the characteristic English contours for expressing **contradiction** and expressing **surprise**.

Continuation rise

The mapping between meaning and tune in English is extremely complex, and linguistic theories of intonation like ToBI have only begun to develop sophisticated models of this mapping. In practice, therefore, most synthesis systems just distinguish two or three tunes, such as the **continuation rise** (at commas), the **question rise** (at question mark if the question is a yes-no question), and a **final fall** otherwise.

8.3.4 More Sophisticated Models: ToBI

While current synthesis systems generally use simple models of prosody like the ones discussed above, recent research focuses on the development of much more sophisticated models. We briefly discuss the **ToBI** and **Tilt** models here.

ToBI

ToBI

One of the most widely used linguistic models of prosody is the **ToBI** (Tone and Break Indices) model (Silverman et al., 1992; Beckman and Hirschberg, 1994; Pierrehumbert, 1980; Pitrelli et al., 1994). ToBI is a phonological theory of intonation that models prominence, tune, and boundaries. ToBI's model of prominence and tunes is based on the five **pitch accents** and four **boundary tones** shown in Fig. 8.9.

Boundary tone

An utterance in ToBI consists of a sequence of intonational phrases, each of which ends in one of the four **boundary tones**. The boundary tones represent the utterance final aspects of tune discussed in Section 8.3.3. Each word in the utterances can optionally be associated with one of the five types of pitch accents.

Each intonational phrase consists of one or more **intermediate phrase**. These phrases can also be marked with kinds of boundary tone, including the **%H** high initial boundary tone, which marks a phrase that is particularly high in the speaker's pitch range, as well as final phrase accents **H-** and **L-**.

Pitch Accents		Boundary Tones	
H*	peak accent	L-L%	“final fall”: “declarative contour” of American English
L*	low accent	L-H%	continuation rise
L*+H	scooped accent	H-H%	“question rise”: cantonal yes-no question contour
L+H*	rising peak accent	H-L%	final level plateau (plateau because H- causes “upstep” of following)
H+!H*	step down		

Figure 8.9 The accent and boundary tones labels from the ToBI transcription system for American English intonation (Beckman and Ayers, 1997; Beckman and Hirschberg, 1994).

Break index

In addition to accents and boundary tones, ToBI distinguishes four levels of phrasing, labeled on a separate **break index** tier. The largest phrasal breaks are the intonational phrase (break index 4) and the intermediate phrase (break index 3), discussed above. Break index 2 is used to mark a disjuncture or pause between words that is smaller than an intermediate phrase, and 1 is used for normal phrase-medial word boundaries.

Tier

Figure 8.10 shows the tone, orthographic, and phrasing tiers of a ToBI transcription, using the Praat program. The same sentence is read with two different tunes. In (a), the word *Marianna* is spoken with a high H* accent, and the sentence has the declarative boundary tone L-L%. In (b), the word *Marianna* is spoken with a low L* accent and the yes-no question boundary tone H-H%. One goal of ToBI is to express different meanings to the different type of accents. Here, the L* accent adds a meaning of *surprise* to the sentence (i.e., with a connotation like ‘Are you really saying it was Marianna?’) (Hirschberg and Pierrehumbert, 1986; Steedman, 2007).

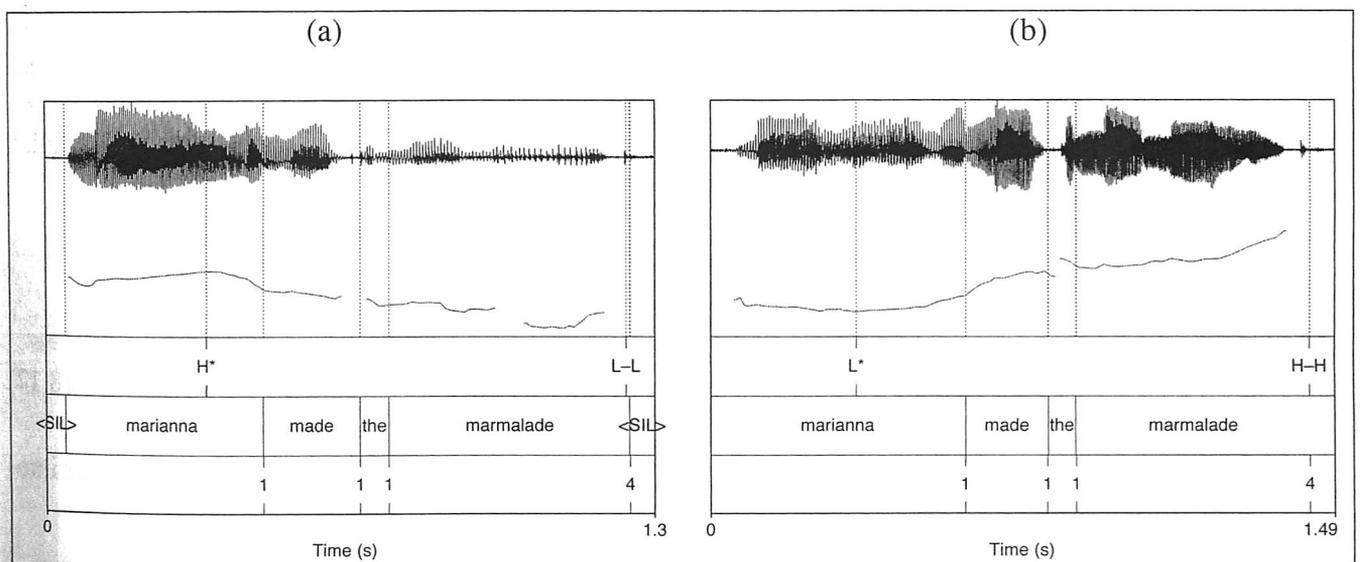


Figure 8.10 The same sentence read by Mary Beckman with two different intonation patterns and transcribed in ToBI. (a) Shows an H* accent and the typical American English declarative final fall L-L%. (b) Shows the L* accent, with the typical American English yes-no question rise H-H%.

ToBI models have been proposed for many languages (Jun, 2005), such as the J_ToBI system for Japanese (Venditti, 2005).

Other Intonation models

Tilt The **Tilt** model (Taylor, 2000) resembles ToBI in that it uses sequences of intonational events like accents and boundary tones. But Tilt does not use ToBI-style discrete phonemic classes for accents. Instead, each event is modeled by continuous parameters that represent the F0 shape of the accent. Instead of giving each event a category label, as in ToBI, each Tilt prosodic event is characterized by a set of three acoustic parameters: the duration, the amplitude, and the **tilt** parameter. These acoustic parameters are trained on a corpus that has been hand-labeled for pitch accents (**a**) and boundary tones (**b**). The human labeling specifies the syllable that bears the accent or tone; the acoustic parameters are then trained automatically from the wavefile. Figure 8.11 shows a sample of a Tilt representation.

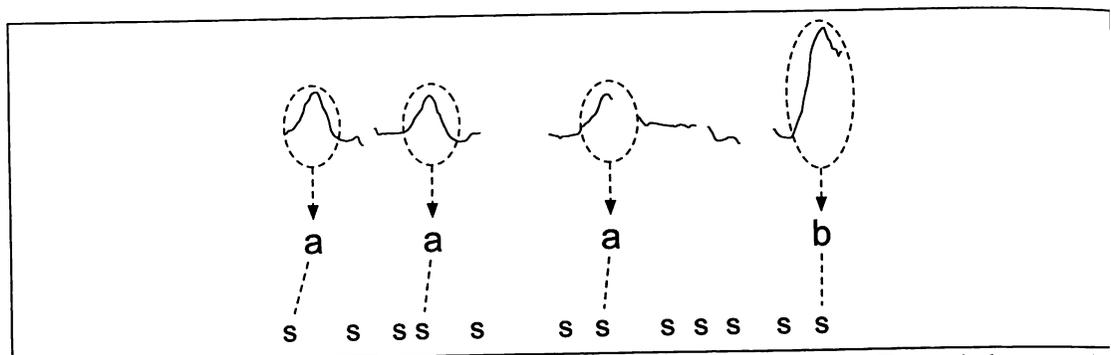


Figure 8.11 Schematic view of events in the Tilt model (Taylor, 2000). Each pitch accent (**a**) and boundary tone (**b**) is aligned with a syllable nucleus *s*.

Each accent in Tilt is viewed as having a (possibly zero) **rise component** up to peak, followed by a (possibly zero) **fall component**. An automatic accent detector finds the start, peak, and end-point of each accent in the wavefile, all of which determine the duration and amplitude of the rise and fall components. The tilt parameter is an abstract description of the F0 slope of an event, calculated by comparison of the relative size of the rise and fall for an event. A tilt value of 1.0 indicates a rise, tilt of -1.0 a fall, 0 equal rise and fall, -0.5 is an accent with a rise and a larger fall, and so on:

$$\begin{aligned} \text{tilt} &= \frac{\text{tilt}_{\text{amp}} + \text{tilt}_{\text{dur}}}{2} \\ &= \frac{|A_{\text{rise}}| - |A_{\text{fall}}|}{|A_{\text{rise}}| + |A_{\text{fall}}|} + \frac{D_{\text{rise}} - D_{\text{fall}}}{D_{\text{rise}} + D_{\text{fall}}} \end{aligned} \quad (8.17)$$

See the Historical Notes section at the end of the chapter for pointers to other intonational models.

8.3.5 Computing Duration from Prosodic Labels

The results of the text analysis processes described so far is a string of phonemes annotated with words with pitch accent marked on relevant words, and appropriate boundary tones marked. For the **unit selection** synthesis approaches that we describe in Section 8.5, this is a sufficient output from the text analysis component.

For **diphone** synthesis, as well as other approaches like formant synthesis, we also need to specify the **duration** and the **F0** values of each segment.

Phones vary quite a bit in duration. Some of the duration is inherent to the identity of the phone itself. Vowels, for example, are generally much longer than consonants; in the Switchboard corpus of telephone speech, the phone [aa] averages 118 milliseconds, and [d] averages 68 milliseconds. But phone duration is also affected by a wide variety of contextual factors, which can be modeled by rule-based or statistical methods.

The most well-known of the rule-based methods is the method of Klatt (1979), which uses rules to model how the average or 'context-neutral' duration of a phone \bar{d} is lengthened or shortened by context, while staying above a minimum duration d_{\min} . Each Klatt rule is associated with a duration multiplicative factor; some examples:

Preasual lengthening:	the vowel or syllabic consonant in the syllable before a pause is lengthened by 1.4.
Non-phrase-final shortening:	segments not phrase-final are shortened by 0.6. Phrase-final postvocalic liquids and nasals are lengthened by 1.4.
Unstressed shortening:	unstressed segments are more compressible, so their minimum duration d_{\min} is halved, and they are shortened by .7 for most phone types.
Lengthening for accent:	a vowel which bears accent is lengthened by 1.4.
Shortening in clusters:	a consonant followed by a consonant is shortened by 0.5.
Pre-voiceless shortening:	vowels are shortened before a voiceless plosive by 0.7.

Given the N factor weights f , the Klatt formula for the duration of a phone is

$$d = d_{\min} + \prod_{i=1}^N f_i \times (\bar{d} - d_{\min}) \quad (8.18)$$

More recent machine learning systems use the Klatt hand-written rules as the basis for defining features, for example, using features such as the following:

- identity of the left and right context phone
- lexical stress and accent values of current phone
- position in syllable, word, phrase
- following pause

We can then train machine learning classifiers like decision trees or the **sum-of-products** model (van Santen, 1994, 1997, 1998), to combine the features to predict the final duration of the segment.

8.3.6 Computing F0 from Prosodic Labels

For diphone, articulatory, HMM, and formant synthesis we also need to specify the F0 values of each segment. For the tone sequence models like ToBI or Tilt, we can do F0 generation in two stages. We first specify F0 **target points** for each pitch accent and boundary tone and then create the F0 contour for the whole sentence by interpolating among these targets (Anderson et al., 1984).

To specify a target point we must first describe what it is (the F0 value) and when it occurs (the exact time at which this peak or trough occurs in the syllable). The F0

Sum-of-products

Target point

Pitch range
 Baseline
 Topline
 Reference line

values of the target points are generally not specified in absolute terms of hertz. Instead, they are defined relative to **pitch range**. A speaker's **pitch range** is the range between the lowest frequency in a particular utterance (the **baseline frequency**) and the highest frequency in the utterance (the **topline**). In some models, target points are specified relative to a line in between called the **reference line**.

For example, we might write a rule specifying that the very beginning of an utterance have a target point of 50% (halfway between the baseline and the topline). In the rule-based system of Jilka et al. (1999), the target point for an H* accent is at 100% (the topline) and for an L* accent at 0% (at the baseline). L+H* accents have two target points, at 20% and 100%. Final boundary tones H-H% and L-L% are extra high and extra low at 120% and -20% respectively.

Second, we must also specify exactly where in the accented syllable the targets apply; this is known as accent **alignment**. In the rule-based system of Jilka et al. (1999), again, H* accents are aligned 60% of the way through the voiced part of the accent syllable (although IP-initial accents are aligned somewhat later in the syllable, and IP-final accents are aligned somewhat earlier).

An alternative to writing these rules by hand is to automatically learn the mapping from pitch accent sequence to F0 value. Black and Hunt (1996), for example, used linear regression to assign target values to each syllable. For each syllable with a pitch accent or boundary tone, they predicted three target values, at the beginning, middle, and end of the syllable. They trained three separate linear regression models, one for each of the three positions in the syllable. Features included the following:

- accent type on the current syllable, two previous and two following syllables
- lexical stress of this syllable and surrounding syllables
- number of syllables to start of phrase and to end of phrase
- number of accented syllables to end of phrase

Such machine learning models require a training set that is labeled for accent; a number of such prosodically labeled corpora exist, although it is not clear how well these models generalize to unseen corpora.

Finally, F0 computation models must model the fact that pitch tends to decline through a sentence; this subtle drop in pitch across an utterance is called **declination**. An example is shown in Fig. 8.12.

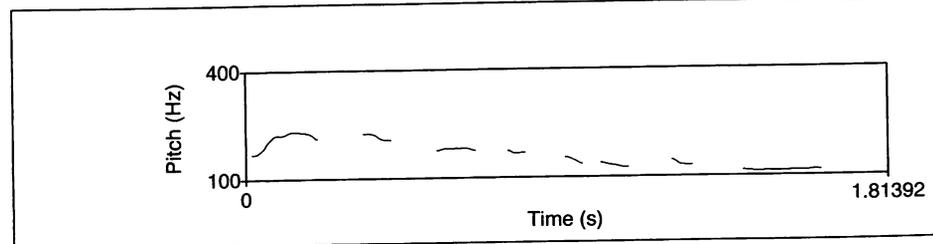


Figure 8.12 F0 declination in the sentence "I was pretty goofy for about 24 hours afterwards"

The exact nature of declination is a subject of much research; some models treat it by allowing the baseline (or both baseline and topline) to decrease slowly over the utterance. In ToBI-like models, this downdrift in F0 is modeled by two separate components; in addition to declination, certain high tones are marked as carrying **downstep**

Downstep

Each downstepped high accent causes the pitch range to be compressed, resulting in a lowered topline for each such accent.

8.3.7 Final Result of Text Analysis: Internal Representation

The final output of text analysis is the **internal representation** of the input text sentence. For unit selection synthesis, the internal representation can be as simple as a phone string together with indications of prosodic boundaries and prominent syllables, as shown in Fig. 8.1 on page 250. For diphone synthesis as well as non-concatenative synthesis algorithms, the internal representation must also include a duration and an F0 value for each phone.

Figure 8.13 shows some sample TTS output from the Festival (Black et al., 1999) diphone speech synthesis system for the sentence “Do you really want to see all of it?”. This output, together with the F0 values shown in Fig. 8.14 would be the input to the **waveform synthesis** component described in Section 8.4. The durations here are computed by a CART-style decision tree (Riley, 1992).

do		you		H* really				want			to	see		L* all		L- H% of it					
d	uw	y	uw	r	ih	l	iy	w	aa	n	t	t	ax	s	iy	ao	l	ah	v	ih	t
110	110	50	50	75	64	57	82	57	50	72	41	43	47	54	130	76	90	44	62	46	220

Figure 8.13 Output of the Festival (Black et al., 1999) generator for the sentence “Do you really want to see all of it?”, together with the F0 contour shown in Fig. 8.14. (Figure thanks to Paul Taylor.)

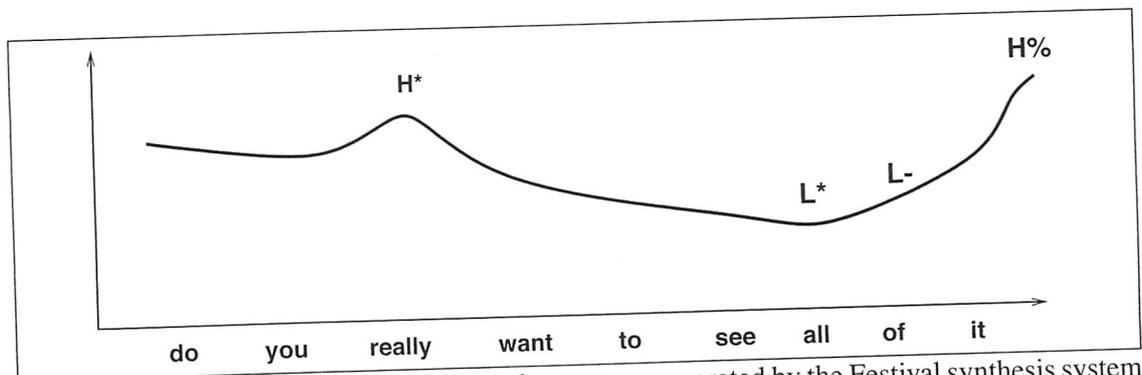


Figure 8.14 The F0 contour for the sample sentence generated by the Festival synthesis system in Fig. 8.13, thanks to Paul Taylor.

As was suggested above, determining the proper prosodic pattern for a sentence is difficult, because we need real-world knowledge and semantic information to know which syllables to accent and which tune to apply. This sort of information is difficult to extract from the text; hence, prosody modules often aim to produce a “neutral declarative” version of the input text and assume the sentence should be spoken in a default way with no reference to discourse history or real-world events. This is one of the main reasons why intonation in TTS often sounds wooden.

8.4 Diphone Waveform Synthesis

We are now ready to see how the internal representation can be turned into a waveform. We present two kinds of **concatenative** synthesis: **diphone synthesis** in this section, and **unit selection synthesis** in the next section.

Recall that for diphone synthesis, our internal representation is as shown in Fig. 8.13 and Fig. 8.14 and consists of a list of phones, each phone associated with a duration and a set of F0 targets.

The diphone concatenative synthesis model generates a waveform from a sequence of phones by selecting and concatenating units from a prerecorded database of **diphones**. A diphone is a phone-like unit going from roughly the middle of one phone to the middle of the following phone. Diphone concatenative synthesis can be characterized by the following steps:

Diphone

Training:

1. Record a single speaker saying an example of each diphone.
2. Cut each diphone from the speech and store all diphones in a database.

Synthesis:

1. Take from the database a sequence of diphones that corresponds to the desired phone sequence.
2. Concatenate the diphones, with slight signal processing at the boundaries.
3. Use signal processing to change the prosody (f0, duration) of the diphone sequence to the desired prosody.

Coarticulation

We tend to use diphones rather than phones for concatenative synthesis because of the phenomenon of **coarticulation**. In Chapter 7 we defined **coarticulation** as the movement of articulators to anticipate the next sound or perseverating movement from the last sound. Because of coarticulation, each phone differs slightly depending on the previous and following phone. Thus if we just concatenated phones we would have very large discontinuities at the boundaries.

In a diphone, we model this coarticulation by including the transition to the next phone inside the unit. The diphone [w-eh], for example, includes the transition from the [w] phone to the [eh] phone. Because a diphone is defined from the middle of one phone to the middle of the next, when we concatenate the diphones, we are concatenating the middle of phones, and the middle of phones tends to be less influenced by the context. Figure 8.15 shows the intuition that the beginning and end of the vowel [eh] have much more movement than the center.

8.4.1 Steps for Building a Diphone Database

Building a diphone database requires six steps:

1. Create a **diphone inventory**.
2. Recruit a speaker.
3. Create a text for the speaker to read for each diphone.

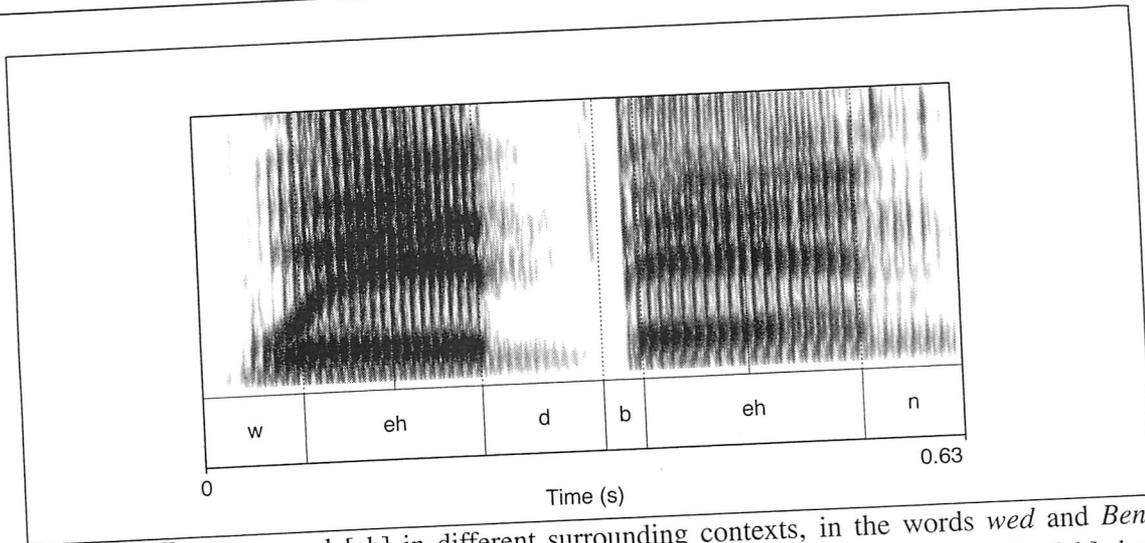


Figure 8.15 The vowel [eh] in different surrounding contexts, in the words *wed* and *Ben*. Notice the differences in the second formants (F2) at the beginning and end of the [eh], but the relatively steady state portion in the middle at the center mark.

4. Record the speaker reading each diphone.
5. Segment, label, and pitch-mark the diphones.
6. Excise the diphones.

What is the inventory of diphones that we need for a system? If we have 43 phones (like the AT&T system of Olive et al. (1998)), there are $43^2 = 1849$ hypothetically possible diphone combinations. Not all of these diphones can actually occur. For example, English **phonotactic** constraints rule out some combinations; phones like [h], [y], and [w] can only occur before vowels. In addition, some diphone systems don't bother storing diphones if there is no possible coarticulation between the phones, such as across the silence between successive voiceless stops. The 43-phone system of Olive et al. (1998) thus has only 1162 diphones rather than the 1849 hypothetically possible set.

Voice talent
Voice

Next, we recruit our speaker, often called a **voice talent**. The database of diphones for this speaker is called a **voice**; commercial systems often have multiple voices, such as one male and one female voice.

Carrier phrase

We'll now create a text for the voice talent to say, and record each diphone. The most important thing in recording diphones is to keep them as consistent as possible; if possible, they should have constant pitch, energy, and duration, so that they are easy to paste together without noticeable breaks. We promote consistency by enclosing each diphone to be recorded in a **carrier phrase**. By surrounding the diphone with other phones, we keep utterance-final lengthening or initial phone effects from making any diphone louder or quieter than the others. We'll need different carrier phrases for consonant-vowel, vowel-consonant, phone-silence, and silence-phone sequences. For example, a consonant vowel sequence like [b aa] or [b ae] could be embedded between the syllables [t aa] and [m aa]:

pause t aa b aa m aa pause
 pause t aa b ae m aa pause
 pause t aa b eh m aa pause

...

If we have an earlier synthesizer voice lying around, we can use that voice to read the prompts out loud and have our voice talent repeat after the prompts. This is another way to keep the pronunciation of each diphone consistent. It is also important to use a high-quality microphone and a quiet room or, better, a sound booth.

Forced alignment

Once we have recorded the speech, we need to label and segment the two phones that make up each diphone, usually by running a speech recognizer in **Forced alignment** mode. In forced alignment mode, a speech recognition is told exactly what the phone sequence is; its job is just to find the exact phone boundaries in the waveform. Speech recognizers are not completely accurate at finding phone boundaries, so usually the automatic phone segmentation is hand-corrected.

We now have the two phones (for example, [b aa]) with hand-corrected boundaries. There are two ways we can create the /b-aa/ diphone for the database. One method is to use rules to decide how far into the phone to place the diphone boundary. For example, for stops, we place the diphone boundary 30% of the way into the phone. For most other phones, we place the diphone boundary 50% into the phone.

Optimal coupling

A more sophisticated way to find diphone boundaries is to store all of both phones and wait to excise the diphones until we are know what phone we are about to concatenate with. In this method, known as **optimal coupling**, we take the two (complete, uncut) diphones we need to concatenate and check every possible cutting point for each diphones, choosing the two cutting points that would make the final frame of the first diphone acoustically most similar to the end frame of the next diphone (Taylor and Isard, 1991; Conkie and Isard, 1996). Acoustic similarity can be measured by **cepstral similarity**, as defined in Section 9.3.

8.4.2 Diphone Concatenation and TD-PSOLA for Prosody

We are now ready to see the remaining steps for synthesizing an individual utterance. Assume that we have completed text analysis for the utterance, have arrived at a sequence of diphones and prosodic targets, and have also grabbed the appropriate sequence of diphones from the diphone database. Next, we need to concatenate the diphones and then adjust the prosody (pitch, energy, and duration) of the diphone sequence to match the prosodic requirements from the intermediate representation.

Click

Given two diphones, what do we need to do to concatenate them successfully? If the waveforms of the two diphones edges across the juncture are very different, a perceptible **click** will result. Thus, we need to apply a windowing function to the edge of both diphones so that the samples at the juncture have low or zero amplitude. Furthermore, if both diphones are voiced, we need to ensure that the two diphones are joined **pitch-synchronously**. This means that the pitch periods at the end of the first diphone must line up with the pitch periods at the beginning of the second diphone; otherwise, the resulting single irregular pitch period at the juncture is perceptible as well.

Pitch-synchronous

TD-PSOLA

Now given our sequence of concatenated diphones, how do we modify the pitch and duration to meet our prosodic requirements? It turns out there is a simple algorithm for doing this, called **TD-PSOLA (Time-Domain Pitch-Synchronous OverLap-and-Add)**.

Pitch marking
Pitch tracking

As we just said, a **pitch-synchronous** algorithm is one in which we do something at each pitch period or **epoch**. For such algorithms it is important to have accurate pitch markings: measurements of exactly where each pitch pulse or **epoch** occurs. An epoch can be defined by the instant of maximum glottal pressure, or alternatively by the instant of glottal closure. Note the distinction between **pitch marking** or **epoch detection** and **pitch tracking**. Pitch tracking gives the value of F_0 (the average cycles per second of the glottis) at each particular point in time, averaged over a neighborhood. Pitch marking finds the exact point in time at each vibratory cycle at which the vocal folds reach some specific point (epoch).

Electroglotto-
graph
EGG
Laryngograph
Lx

Epoch labeling can be done in two ways. The traditional way, and still the most accurate, is to use an **electroglottograph** or **EGG** (often also called a **laryngograph** or **Lx**). An EGG is a device that straps onto the (outside of the) speaker's neck near the larynx and sends a small current through the Adam's apple. A transducer detects whether the glottis is open or closed by measuring the impedance across the vocal folds. Some modern synthesis databases are still recorded with an EGG. The problem with using an EGG is that it must be attached to the speaker while he or she is recording the database. Although an EGG isn't particularly invasive, it is still annoying, and the EGG must be used during recording; it can't be used to pitch-mark speech that has already been collected. Modern epoch detectors are now approaching such a level of accuracy that EGGs are no longer used in most commercial TTS engines. Algorithms for epoch detection include Brookes and Loke (1999), and Veldhuis (2000).

Overlap-and-add
OLA

Given an epoch-labeled corpus, the intuition of TD-PSOLA is that we can modify the pitch and duration of a waveform by extracting a frame for each pitch period (windowed so that the frame doesn't have sharp edges) and then recombining these frames in various ways by simply overlapping and adding the windowed pitch period frames (we introduce the idea of windows in Section 9.3.2). The idea that we modify a signal by extracting frames, manipulating them in some way, and then recombining them by adding up the overlapped signals is called the **overlap-and-add** or **OLA** algorithm; TD-PSOLA is a special case of overlap-and-add in which the frames are pitch synchronous and the whole process takes place in the time domain.

For example, to assign a specific duration to a diphone, we might want to lengthen the recorded master diphone. To lengthen a signal with TD-PSOLA, we simply insert extra copies of some of the pitch-synchronous frames, essentially duplicating a piece of the signal. Figure 8.16 shows the intuition.

TD-PSOLA can also be used to change the F_0 value of a recorded diphone to give a higher or lower value. To increase the F_0 , we extract each pitch-synchronous frame from the original recorded diphone signal, place the frames closer together (overlapping them), with the amount of overlap determined by the desired period and hence frequency, and then add up the overlapping signals to produce the final signal. But note that by moving all the frames closer together, we make the signal shorter in time! Thus to change the pitch while holding the duration constant, we need to add duplicate frames.

Figure 8.17 shows the intuition; in this figure we have explicitly shown the extracted pitch-synchronous frames that are overlapped and added; note that the frames moved closer together (increasing the pitch) while extra frames have been added to hold the duration constant.

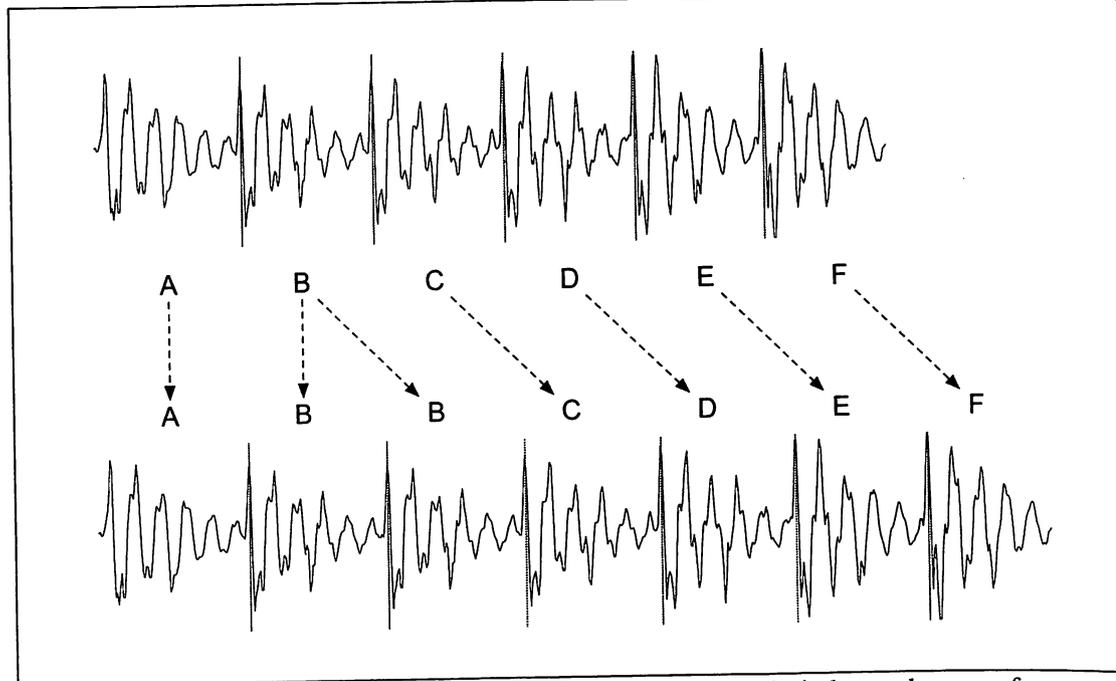


Figure 8.16 TD-PSOLA for duration modification. Individual pitch-synchronous frames can be duplicated to lengthen the signal (as shown here), or deleted to shorten the signal.

8.5 Unit Selection (Waveform) Synthesis

Diphone waveform synthesis suffers from two main problems. First, the stored diphone database must be modified by signal process methods like PSOLA to produce the desired prosody. Any kind of signal processing of the stored speech leaves artifacts in the speech that can make the speech sound unnatural. Second, diphone synthesis captures only the coarticulation due to a single neighboring phone. But there are many more global effects on phonetic realization, including more distant phones, syllable structure, the stress patterns of nearby phones, and even word-level effects.

*Unit selection
synthesis*

For this reason, modern commercial synthesizers are based on a generalization of diphone synthesis called **unit selection synthesis**. Like diphone synthesis, unit selection synthesis is a kind of concatenative synthesis algorithm. The word **unit** means any stored piece of speech that is concatenated together to form an output. The intuition of unit selection synthesis is that we can store units of different sizes, which can be much larger than diphones. Unit selection thus differs from classic diphone synthesis in two ways:

1. In diphone synthesis, the database stores one copy of each diphone; in unit selection, the database is many hours long, with many copies of each diphone.
2. In diphone synthesis, the prosody of the units is modified by PSOLA or similar algorithms; in unit selection no (or minimal) signal processing is applied to the concatenated units.

The strengths of unit selection are due to the large unit database. In a sufficiently large database, entire words or phrases of the utterance we want to synthesize may already be present in the database, resulting in an extremely natural waveform for these

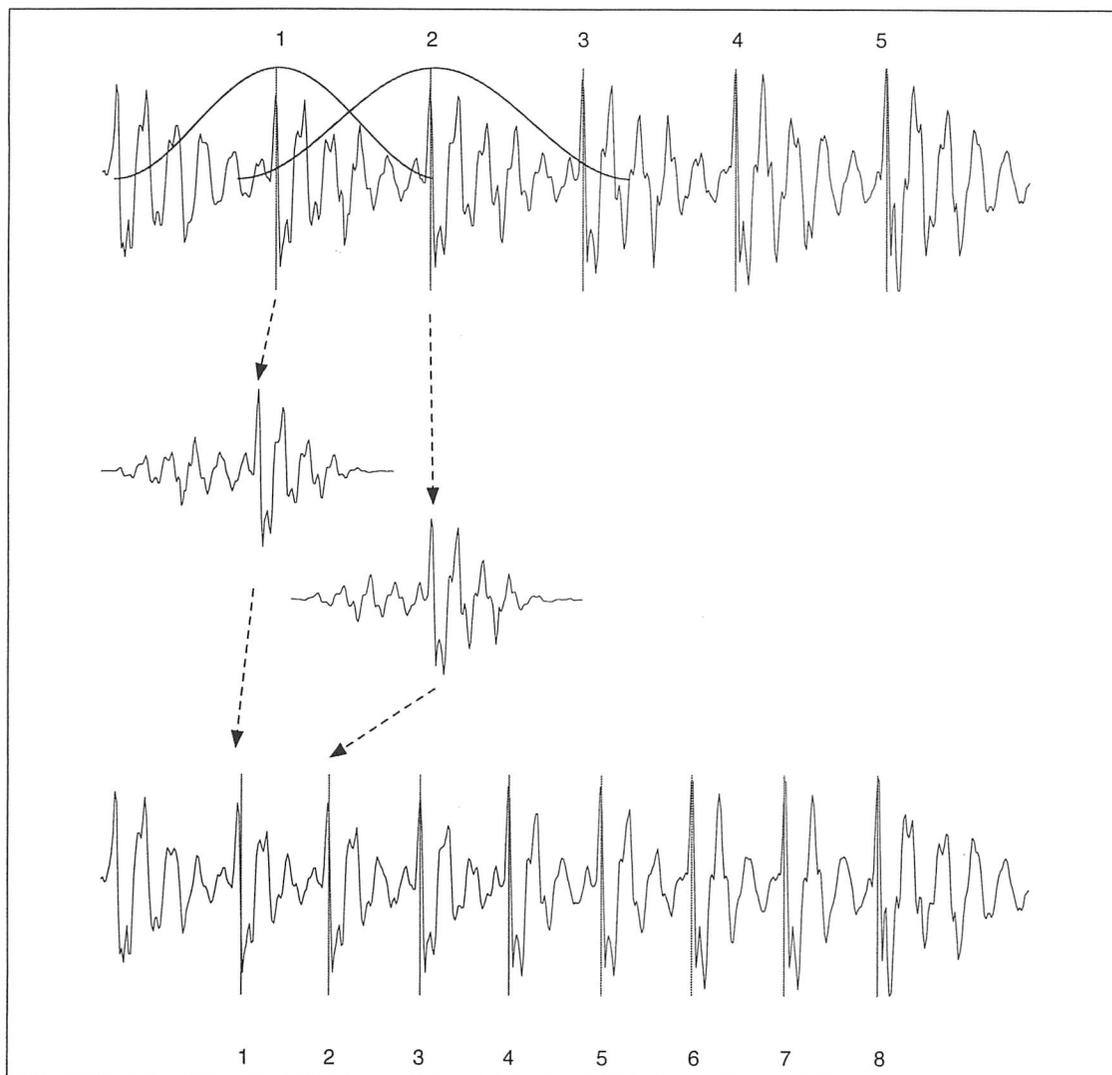


Figure 8.17 TD-PSOLA for pitch (F_0) modification. To increase the pitch, we extract the individual pitch-synchronous frames, Hanning-window them, move them closer together and then add them up. To decrease the pitch, we move the frames further apart. Increasing the pitch will result in a shorter signal (since the frames are closer together), so we also need to duplicate frames if we want to change the pitch while holding the duration constant.

words or phrases. Thus we implicitly create larger units by selecting diphones that are consecutive in the database. In addition, in cases in which we can't find a large chunk and have to back off to individual diphones, many copies of each diphone make it more likely that we will find one that will fit in naturally.

The architecture of unit selection can be summarized as follows. We are given a large database of units; let's assume these are diphones (although it's also possible to do unit selection with other kinds of units such as half-phones, syllables, or half-syllables). We are also given a characterization of the target 'internal representation', that is, a phone string together with features such as stress values, word identity, F_0 information, as described in Fig. 8.1 on page 250.

The goal of the synthesizer is to select from the database the best sequence of diphone units that corresponds to the target representation. What do we mean by the "best" sequence? Intuitively, the best sequence would be one in which

- each diphone unit we select exactly meets the specifications of the target diphone (in terms of F0, stress level, phonetic neighbors, etc.)
- each diphone unit concatenates smoothly with its neighboring units, with no perceptible break

Of course, in practice, we can't guarantee that there will be a unit that exactly meets our specifications, and we are unlikely to find a sequence of units in which every single join is imperceptible. Thus, in practice, unit selection algorithms implement a gradient version of these constraints, and attempt to find the sequence of unit that at least minimizes the **target cost** and the **join cost**:

Target cost

Join cost

Target cost $T(u_t, s_t)$: how well the target specification s_t matches the potential unit u_t

Join cost $J(u_t, u_{t+1})$: how well (perceptually) the potential unit u_t joins with its potential neighbor u_{t+1}

The T and J values are expressed as **costs**, meaning that high values indicate bad matches and bad joins (Hunt and Black, 1996).

Formally, then, the task of unit selection synthesis, given a sequence S of T target specifications, is to find the sequence \hat{U} of T units from the database that minimizes the sum of these costs:

$$\hat{U} = \operatorname{argmin}_U \sum_{t=1}^T T(s_t, u_t) + \sum_{t=1}^{T-1} J(u_t, u_{t+1}) \quad (8.19)$$

Let's first define the target cost and the join cost in more detail before we turn to the decoding and training tasks.

The target cost measures how well the unit matches the target diphone specification. We can think of the specification for each diphone target as a feature vector; here are three sample vectors for three target diphone specifications, using dimensions (features) like *should the syllable be stressed*, and *where in the intonational phrase should the diphone come from*:

```
/ih-t/, +stress, phrase internal, high F0, content word
/n-t/, -stress, phrase final, high F0, function word
/dh-ax/, -stress, phrase initial, low F0, word 'the'
```

We'd like the distance between the target specification s and the unit to be some function of how different the unit is on each of these dimensions from the specification. Let's assume that for each dimension p , we can come up with some **subcost** $T_p(s_t[p], u_j[p])$. The subcost for a binary feature like *stress* might be 1 or 0. The subcost for a continuous feature like F0 might be the difference (or log difference) between the specification F0 and unit F0. Since some dimensions are more important to speech perceptions than others, we'll also want to weight each dimension. The simplest way to combine all these subcosts is just to assume that they are independent and additive. Using this model, the total target cost for a given target/unit pair is the weighted sum over all these subcosts for each feature/dimension:

$$T(s_t, u_j) = \sum_{p=1}^P w_p T_p(s_t[p], u_j[p]) \quad (8.20)$$

The target cost is a function of the desired diphone specification and a unit from the database. The **join cost**, by contrast, is a function of two units from the database. The goal of the join cost is to be low (0) when the join is completely natural, and high when the join would be perceptible or jarring. We meet this goal by measuring the acoustic similarity of the edges of the two units that we will be joining. If the two units have very similar energy, F0, and spectral features, they will probably join well. Thus as with the target cost, we compute a join cost by summing weighted subcosts:

$$J(u_t, u_{t+1}) = \sum_{p=1}^P w_p J_p(u_t[p], u_{t+1}[p]) \quad (8.21)$$

The three subcosts used in the classic Hunt and Black (1996) algorithm are the **cepstral distance** at the point of concatenation, and the absolute differences in both log power and F0. We introduce the cepstrum in Section 9.3.

In addition, if the two units u_t and u_{t+1} to be concatenated were consecutive diphones in the unit database (i.e., they followed each other in the original utterance), then we set the join cost to 0: $J(u_t, u_{t+1}) = 0$. This is an important feature of unit selection synthesis since it encourages large natural sequences of units to be selected from the database.

How do we find the best sequence of units that minimizes the sum of the target and join costs as expressed in Eq. 8.19? The standard method is to think of the unit selection problem as a hidden Markov model. The target units are the observed outputs, and the units in the database are the hidden states. Our job is to find the best hidden state sequence. We can use the Viterbi algorithm (Chapter 6) to solve this problem. Figure 8.18 shows a sketch of the search space as well as the best (Viterbi) path that determines the best unit sequence.

The weights for join and target costs are often set by hand because the number of weights is small (on the order of 20) and machine learning algorithms don't always achieve human performance. The system designer listens to entire sentences produced by the system and chooses values for weights that result in reasonable sounding utterances. Various automatic weight-setting algorithms do exist, however. Many of these assume we have some sort of distance function between the acoustics of two sentences, perhaps based on cepstral distance. The method of Hunt and Black (1996), for example, holds out a test set of sentences from the unit selection database. For each of these test sentences, we synthesize from the word sequence a sentence waveform (using units from the other sentences in the training database). Next we compare the acoustics of the synthesized sentence with the acoustics of the true human sentence. We now have a sequence of synthesized sentences, each one associated with a distance function to its human counterpart. We then use linear regression based on these distances to set the target cost weights so as to minimize the distance.

There are also more advanced methods of assigning both target and join costs. For example, above we computed target costs between two units by looking at the features

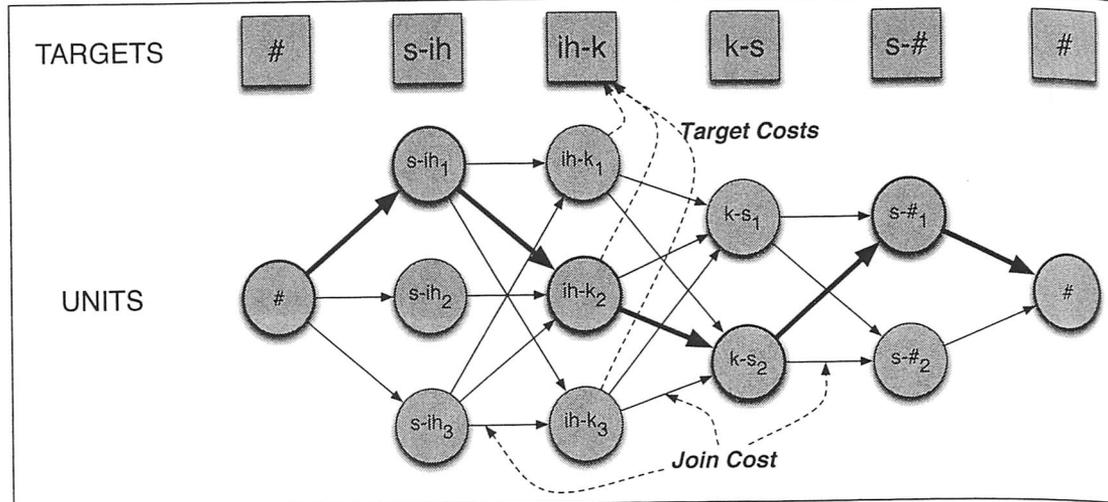


Figure 8.18 Decoding in unit selection, showing the sequence of target (specification) diphones for the word *six* and the set of possible database diphone units that we must search through. The best (Viterbi) path that minimizes the sum of the target and join costs is in bold.

of the two units, doing a weighted sum of feature costs, and choosing the lowest-cost unit. An alternative approach (which the new reader might need to come back to after learning the speech recognition techniques introduced in the next chapters) is to map the target unit into some acoustic space and then find a unit that is near the target in that acoustic space. The method of Donovan and Woodland (1995) and Donovan and Eide (1998), for example, clusters all the training units by using the decision tree algorithm of speech recognition described in Fig. 10.14 of Section 10.3. The decision tree is based on the same features described above, but here for each set of features, we follow a path down the decision tree to the leaf node that contains a cluster of units that have those features. This cluster of units can be parameterized by a Gaussian model just as for speech recognition, so that we can map a set of features into a probability distribution over cepstral values and hence easily compute a distance between the target and a unit in the database. As for join costs, more sophisticated metrics make use of how perceivable a particular join might be (Wouters and Macon, 1998; Syrdal and Conkie, 2004; Bulyko and Ostendorf, 2001).

8.6 Evaluation

Speech synthesis systems are evaluated by human listeners. The development of a good automatic metric for synthesis evaluation, one that would eliminate the need for expensive and time-consuming human listening experiments, remains an open and exciting research topic.

Intelligibility

Quality

The minimal evaluation metric for speech synthesis systems is **intelligibility**: the ability of a human listener to correctly interpret the words and meaning of the synthesized utterance. A further metric is **quality**: an abstract measure of the naturalness, fluency, or clarity of the speech.

Diagnostic Rhyme
Test
DRT

The most local measures of intelligibility test the ability of a listener to discriminate between two phones. The **Diagnostic Rhyme Test (DRT)** (Voiers et al., 1975) tests the intelligibility of initial consonants. It is based on 96 pairs of confusable rhyming words differing only in a single phonetic feature, such as *dense/tense* or *bond/pond* (differing in voicing) or *mean/beat* or *neck/deck* (differing in nasality). For each pair, listeners hear one member of the pair and indicate which they think it is. The percentage of right answers is then used as an intelligibility metric. The **Modified Rhyme Test (MRT)** (House et al., 1965) is a similar test based on a different set of 300 words, consisting of 50 sets of 6 words. Each 6-word set differs in either initial or final consonants (e.g., *went, sent, bent, dent, tent, rent* or *bat, bad, back, bass, ban, bath*). Listeners are given a single word they must identify from a closed list of six words; the percentage of correct identifications is again used as an intelligibility metric.

Modified Rhyme
Test
MRT

Carrier phrase

Since context effects are very important, both DRT and MRT words are embedded in **carrier phrases** like the following:

Now we will say <word> again.

SUS

To test units larger than single phones, we can use **semantically unpredictable sentences (SUS)** (Benoît et al., 1996). We construct such sentences by taking a simple POS template like DET ADJ NOUN VERB DET NOUN and inserting random English words in the slots, to produce sentences like

The unsure steaks closed the fish.

Measures of intelligibility like DRT/MRT and SUS factor out the role of context in measuring intelligibility. While this allows for a carefully controlled measure of a system's intelligibility, such acontextual or semantically unpredictable sentences aren't a good fit to how TTS is used in commercial applications. Thus, in commercial applications instead of DRT or SUS, we generally test intelligibility in situations that mimic the desired applications: reading addresses out loud, reading lines of news text, and so on.

MOS

To further evaluate the **quality** of the synthesized utterances, we can play a sentence for listeners and ask them to give a **mean opinion score (MOS)**, a rating of how good the synthesized utterances are, usually on a scale from 1–5. We can then compare systems by comparing their MOS scores on the same sentences (using, e.g., t-tests to test for significant differences).

AB tests

If we are comparing exactly two systems (perhaps to see if a particular change actually improved the system), we can use **AB tests**. In AB tests, we play the same sentence synthesized by two different systems (an A and a B system). The human listeners choose which of the two utterances they like better. We do this for 50 sentences and compare the number of sentences preferred for each system. To avoid listener bias, for each sentence we must present the two synthesized waveforms in random order.

Bibliographical and Historical Notes

As we noted at the beginning of the chapter, speech synthesis is one of the earliest fields of speech and language processing. The 18th century saw a number of physical models

of the articulation process, including the von Kempelen model mentioned above, as well as the 1773 vowel model of Kratzenstein in Copenhagen using organ pipes.

But the modern era of speech synthesis can clearly be said to have arrived by the early 1950s, when all three of the major paradigms of waveform synthesis—formant synthesis, articulatory synthesis, and concatenative synthesis—had been proposed.

Concatenative synthesis seems to have been first proposed by Harris (1953) at Bell Laboratories; he literally spliced together pieces of magnetic tape corresponding to phones. Harris's proposal was actually more like unit selection synthesis than diphone synthesis, in that he proposed storing multiple copies of each phone and proposed the use of a join cost (choosing the unit with the smoothest formant transitions with the neighboring unit). Harris's model was based on the phone, rather than the diphone, resulting in problems due to coarticulation. Peterson et al. (1958) added many of the basic ideas of unit selection synthesis, including the use of diphones, a database with multiple copies of each diphone with differing prosody, each labeled with prosodic features including F0, stress, and duration, and the use of join costs based on F0 and formant distance between neighboring units. They also proposed microconcatenation techniques such as windowing the waveforms. The Peterson et al. model was purely theoretical, however, and concatenative synthesis was not implemented until the 1960s and 1970s, when diphone synthesis was first implemented (Dixon and Maxey, 1968; Olive, 1977). Later diphone systems included larger units such as consonant clusters (Olive and Liberman, 1979). Modern unit selection, including the idea of large units of non-uniform length and the use of a target cost, was invented by Sagisaka (1988). Sagisaka et al. (1992). Hunt and Black (1996) formalized the model and put it in the form in which we have presented it in this chapter in the context of the ATR CHATR system (Black and Taylor, 1994). The idea of automatically generating synthesis units by clustering was first invented by Nakajima and Hamada (1988), but was developed mainly by (Donovan, 1996) who incorporated decision tree clustering algorithms from speech recognition. Many unit selection innovations took place as part of the AT&T NextGen synthesizer (Syrdal et al., 2000; Syrdal and Conkie, 2004).

We focused in this chapter on concatenative synthesis, but there are two other synthesis paradigms: **formant synthesis**, in which we try to build rules that generate artificial spectra, including especially formants; and **articulatory synthesis**, in which we try to directly model the physics of the vocal tract and articulatory process.

Formant synthesizers originally were inspired by attempts to mimic human speech by generating artificial spectrograms. The Haskins Laboratories Pattern Playback Machine generated a sound wave by painting spectrogram patterns on a moving transparent belt and using reflectance to filter the harmonics of a waveform (Cooper et al. 1951); other very early formant synthesizers include those of Lawrence (1953) and Farabee (1951). Perhaps the most well-known of the formant synthesizers were the **Klatt formant synthesizer** and its successor systems, including the MITalk system (Allen et al. 1987) and the Klattalk software used in Digital Equipment Corporation's DECtalk (Klatt, 1982). See Klatt (1975) for details.

Articulatory synthesizers attempt to synthesize speech by modeling the physics of the vocal tract as an open tube. Representative models, both early and somewhat more recent include those of Stevens et al. (1953), Flanagan et al. (1975), and Farabee (1986). See Klatt (1975) and Flanagan (1972) for more details.

Development of the text analysis components of TTS came somewhat later, as techniques were borrowed from other areas of natural language processing. The input to early synthesis systems was not text, but rather phonemes (typed in on punched cards). The first text-to-speech system to take text as input seems to have been the system of Umeda and Teranishi (Umeda et al., 1968; Teranishi and Umeda, 1968; Umeda, 1976). The system included a lexicalized parser that assigned prosodic boundaries, as well as accent and stress; the extensions in Coker et al. (1973) added more rules, for example, for deaccenting light verbs, and explored articulatory models as well. These early TTS systems used a pronunciation dictionary for word pronunciations. To expand to larger vocabularies, early formant-based TTS systems such as MITalk (Allen et al., 1987) used letter-to-sound rules instead of a dictionary since computer memory was far too expensive to store large dictionaries.

Modern grapheme-to-phoneme models derive from the influential early probabilistic grapheme-to-phoneme model of Lucassen and Mercer (1984), which was originally proposed in the context of speech recognition. The widespread use of such machine learning models was delayed, however, because early anecdotal evidence suggested that hand-written rules worked better than, for example, the neural networks of Sejnowski and Rosenberg (1987). The careful comparisons of Damper et al. (1999) showed that machine learning methods were generally superior. A number of such models make use of pronunciation by analogy (Byrd and Chodorow, 1985; Dedina and Nusbaum, 1991; Daelemans and van den Bosch, 1997; Marchand and Damper, 2000) or latent analogy (Bellegarda, 2005); HMMs (Taylor, 2005) have also been proposed. The most recent work makes use of joint **grapheme** models, in which the hidden variables are phoneme-grapheme pairs and the probabilistic model is based on joint rather than conditional likelihood (Deligne et al., 1995; Luk and Damper, 1996; Galescu and Allen, 2001; Bisani and Ney, 2002; Chen, 2003).

The literature on prosody is vast; another important computational model, for example, is the **Fujisaki** model (Fujisaki and Ohno, 1997). IViE (Grabe, 2001) is an extension of ToBI that focuses on labeling different varieties of English (Grabe et al., 2000). There is also much debate on the units of intonational structure, including **intonational phrases** (Beckman and Pierrehumbert, 1986), **intonation units** (Du Bois et al., 1983) or **tone units** (Crystal, 1969), and their relation to clauses and other syntactic units (Chomsky and Halle, 1968; Langendoen, 1975; Streeter, 1978; Hirschberg and Pierrehumbert, 1986; Selkirk, 1986; Nespor and Vogel, 1986; Croft, 1995; Ladd, 1996; Ford and Thompson, 1996; Ford et al., 1996). Much recent work on speech synthesis has focused on generating emotional speech (Cahn, 1990; Bulut et al., 2002; Hamza et al., 2004; Eide et al., 2004; Lee et al., 2006; Schroder, 2006, among others).

One of the most exciting new paradigms for speech synthesis is **HMM synthesis**, first proposed by Tokuda et al. (1995b) and elaborated in Tokuda et al. (1995a), Tokuda et al. (2000), and Tokuda et al. (2003). See also the textbook summary of HMM synthesis in Taylor (2008). Huang et al. (2001) and Gibbon et al. (2000) present more information on TTS evaluation. See also the annual speech synthesis competition called the **Blizzard Challenge** (Black and Tokuda, 2005; Bennett, 2005).

Two classic text-to-speech synthesis systems are described in Allen et al. (1987) (the MITalk system) and Sproat (1998b) (the Bell Labs system). Recent textbooks include Dutoit (1997), Huang et al. (2001), Taylor (2008), and Alan Black's online lec-

Grapheme

Fujisaki

Intonation unit

Tone unit

HMM synthesis

Blizzard
Challenge

ture notes at http://festvox.org/festtut/notes/festtut_toc.html. Influential collections of papers include van Santen et al. (1997), Sagisaka et al. (1997), Narayanan and Alwan (2004). Conference publications appear in the main speech engineering conferences (INTERSPEECH, *IEEE ICASSP*) and the *Speech Synthesis Workshops*. Journals include *Speech Communication*, *Computer Speech and Language*, the *IEEE Transactions on Audio, Speech, and Language Processing*, and the *ACM Transactions on Speech and Language Processing*.

Exercises

- 8.1 Implement the text normalization routine that deals with MONEY, that is, mapping strings of dollar amounts like \$45, \$320, and \$4100 to words (either writing code directly or designing an FST). If there are multiple ways to pronounce a number you may pick your favorite way.
- 8.2 Implement the text normalization routine that deals with NTEL, that is, seven-digit phone numbers like 555-1212, 555-1300, and so on. Use a combination of the **paired** and **trailing unit** methods of pronunciation for the last four digits. (Again, either write code or design an FST).
- 8.3 Implement the text normalization routine that deals with type NDATE in Fig. 8.4.
- 8.4 Implement the text normalization routine that deals with type NTIME in Fig. 8.4.
- 8.5 (Suggested by Alan Black.) Download the free Festival speech synthesizer. Augment the lexicon to correctly pronounce the names of everyone in your class.
- 8.6 Download the Festival synthesizer. Using your own voice, record and train a diphone synthesizer.
- 8.7 Build a phrase boundary predictor. You can use any classifier you like, and you should implement some of the features described on page 263.