Natural Language Engineering **2** (4) 369–380. Printed in the United Kingdom

© 1996 Cambridge University Press

Multilingual text analysis for text-to-speech synthesis

RICHARD SPROAT

Language Modeling Research Department, Bell Laboratories, Lucent Technologies, 700 Mountain Avenue, Room 2d-451, Murray Hill, NJ, 07974–0636, USA
e-mail: rws@bell-labs.com

(*Received 28 December 1996; Revised 6 February 1997*)

Abstract

We present a model of text analysis for text-to-speech (TTS) synthesis based on (weighted) finite state transducers, which serves as the text analysis module of the multilingual Bell Labs TTS system. The transducers are constructed using a lexical toolkit that allows declarative descriptions of lexicons, morphological rules, numeral-expansion rules, and phonological rules, inter alia. To date, the model has been applied to eight languages: Spanish, Italian, Romanian, French, German, Russian, Mandarin and Japanese.

1 Introduction

The first task faced by any text-to-speech (TTS) system is the conversion of input text into an internal linguistic representation. This is in general a complex task since the written form of any language is at best an imperfect representation of the corresponding spoken forms. Among the problems that one faces in handling ordinary text are the following:

1. While a large number of languages delimit words using whitespace or some other device, some languages, such as Chinese and Japanese do not. One is therefore required to 'reconstruct' word boundaries in TTS systems for these languages.
2. Digit sequences need to be expanded into words, and more generally into well-formed number names: so *243* in English would generally be expanded as *two hundred and forty three*.
3. Abbreviations need to be expanded into full words. In general this can involve some amount of contextual disambiguation: so *kg*. can be either *kilogram* or *kilograms*, depending upon the context.
4. Ordinary words and names need to pronounced. In many languages, this requires morphological analysis: even in languages with fairly 'regular' spelling, morphological structure is often crucial in determining the pronunciation of a word.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

5. Prosodic phrasing is only sporadically indicated (by punctuation marks) in the input, and phrasal accentuation is almost never indicated. At a minimum, some amount of lexical analyis (to determine, e.g. grammatical part of speech) is necessary to predict which words to make prominent, and where to place prosodic boundaries.

In many TTS systems the first three tasks — word segmentation, and digit and abbreviation expansion — would be classed under the rubric of *text normalization* and would generally be handled prior to, and often in a quite differerent fashion from the last two problems, which fall more squarely within the domain of linguistic analysis[1].

One problem with this approach is that in many cases the selection of the correct linguistic form for a 'normalized' item cannot be chosen before one has done a certain amount of linguistic analysis. Consider an example that is problematic for the Bell Laboratories American English TTS system, a system that treats text normalization prior to, and separately from, the rest of linguistic analysis. If one encounters the string *$5* in an English text, the normal expansion would be *five dollars*. But this expansion is not always correct: when functioning as a prenominal modifier, as in the phrase *$5 bill*, the correct expansion is *five dollar*, since in general plural noun forms cannot function as modifiers in English. The analysis of complex noun phrases in the American English system (Sproat 1994) comes later than the preprocessing phase, and since a hard decision has been made in the earlier phase, the system produces an incorrect result.

An even more compelling example can be found in Russian. While in English the percentage symbol '%', when denoting a percentage, is always read as *percent*, in Russian selecting the correct form depends on complex contextual factors. The first decision that needs to be made is whether or not the number-percent string is modifying a following noun. Russian in general disallows noun-noun modification: in constructions equivalent to noun-noun compounds in English, the first noun must be converted into an adjective: thus *rog* 'rye', but *rzhanoj xleb* (rye+adj bread) 'rye bread'. This general constraint applies equally to *procent* 'percent', so that the correct rendition of *20% skidka* 'twenty percent discount' is *dvadcati-procentnaja skidka* (twenty$_{[gen]}$-percent+adj$_{[nom,sg,fem]}$ discount$_{[nom,sg,fem]}$). Note that not only does *procent* have to be in the adjectival form, but as with any Russian adjective it must also agree in number, case and gender with the following noun. Observe also that the word for 'twenty' must occur in the genitive case. In general, numbers which modify adjectives in Russian must occur in the genitive case: consider, for example, *etazh* 'storey', and *dvux-etazhnyj* (two$_{[gen]}$-storey+adj$_{[nom,sg,masc]}$). If the percentage expression is not modifying a following noun, then the nominal form *procent* is used. However this form appears in different cases depending upon the number it occurs with. With numbers ending in *one* (including compound numbers like *twenty one*), *procent* occurs in the nominative singular. After so-called paucal numbers

[1] But see Traber (1995), who treats numeral expansion as an instance of morphological analysis, and also the work of van Leeuwen (1989), who uses cascaded rewrite rules within his TooLiP system towards the same ends.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

— *two*, *three*, *four* and their compounds — the genitive singular *procenta* is used. After all other numbers one finds the genitive plural *procentov*. So we have *odin procent* (one percent$_{[nom,sg]}$), *dva procenta* (two percent$_{[gen,sg]}$), and *pyat' procentov* (five percent$_{[gen,pl]}$). All of this, however, presumes that the percentage expression as a whole is in a non-oblique case. If the expression is in an oblique case, then both the number and *procent* show up in that case, with *procent* being in the singular if the number ends in *one*, and the plural otherwise: *s odnym procentom* (with one$_{[instr,sg,masc]}$ percent$_{[instr,sg]}$) 'with one percent'; *s pjat'ju procentami* (with five$_{[instr,pl]}$ percent$_{[instr,pl]}$) 'with five percent'. As with the adjectival forms, there is nothing peculiar about the behavior of the noun *procent*: all nouns exhibit similar behavior in combination with numbers (Franks 1994). The complexity, of course, arises because the written form '*%*' gives no indication of what linguistic form it corresponds to. Furthermore, there is no way to correctly expand this form without doing a substantial amount of analysis of the context, including some analysis of the morphological properties of the surrounding words, as well as an analysis of the relationship of the percentage expression to those words[2].

The obvious solution to this problem is to delay the decision on how exactly to transduce symbols like '$' in English or '%' in Russian until one has enough information to make the decision in an informed manner. This suggests a model where, say, an expression like '20%' in Russian is transduced into all possible renditions, and the correct form selected from the lattice of possibilities by filtering out the illegal forms[3]. An obvious computational mechanism for accomplishing this is the *finite state transducer* (FST). Indeed, since it is well known that FSTs can also be used to model (most) morphology and phonology (Koskenniemi 1983; Karttunen *et al.* 1992; Sproat 1992), can segment words in Chinese text and perform other tokenization tasks (Sproat *et al.* 1996; Karttunen *et al.* 1997), and, as we show here, can perform text analysis operations such as numeral expansion, this suggests a model of text analysis that is entirely based on regular relations. We present such a model below. More specifically we present a model of text analysis for TTS based on *weighted* FSTs (WFSTs) (Pereira *et al.* 1994; Pereira and Riley 1996), which serves as the text analysis module of the multilingual Bell Labs TTS system[4]. To date, the model has been applied to eight languages: Spanish, Italian, Romanian, French, German, Russian, Mandarin and Japanese. One property of this model which distinguishes it from most text analyzers used in TTS systems is that there is no sense in which such tasks as numeral expansion or word segmentation

[2] As Kimmo Koskenniemi points out, other languages also present similar complexities. In Finnish the phrase 'to 521 pupils' might be written *521 oppilaalle* (inessive case). This would be rendered as *viide+lle+sada+lle+kahde+lle+kymmene+lle+yhde+lle oppilaa+lle* (five + INESSIVE + hundred + INESSIVE + two + INESSIVE + ten + INESSIVE + one + INESSIVE pupil + INESSIVE).

[3] The model is thus similar in spirit to work on constraint grammar: see, for example, (Voutilainen 1994).

[4] For lack of space we cannot elaborate on the algorithmic differences between the use of weighted transducers and more well-known unweighted transducers. The reader is referred to (Pereira and Riley 1996) for discussion and further references.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

are logically prior to other aspects of linguistic analysis, and there is therefore no distinguished 'text normalization' phase.

2 Overall architecture

Let us start with the example of the lexical analysis and pronunciation of ordinary words, taking again an example from Russian. Russian orthography is often described as 'morphological' (Daniels and Bright 1996), meaning that the orthography represents not a surface phonemic level of representation, but a more abstract level. This description is correct, but from the point of view of predicting word pronunciation, it is noteworthy that Russian, with a well-defined set of lexical exceptions, is almost completely phonemic in that one can predict the pronunciation of most words in Russian based on the spelling of those words — provided one knows the placement of word stress, since several Russian vowels undergo reduction to varying degrees depending upon their position relative to stressed syllables. The catch is that word stress usually depends upon knowing lexical properties of the word, including morphological class information. To take a concrete example, consider the word *kostra* (Cyrillic кострa) (bonfire$_{[gen,sg]}$). This word belongs to a class of masculine nouns where the word stress is placed on the inflectional ending, where there is one. Thus the stress pattern is *kostr′a*, and the pronunciation is /kəstr′ʌ/, with the first /o/ reduced to /ə/. Let us assume that the morphological representation for this word is something like *kostËr*{*noun*}{*masc*}{*inan*}+′*a*{*sg*}{*gen*}, where for convenience we represent phonological and morphosyntactic information as part of the same string[5]. Assuming a finite state model of lexical structure (Koskenniemi 1983; Karttunen *et al.* 1992), we can easily imagine a set of transducers M that map from that level into a level that gives what we shall term the *minimal morphologically-motivated annotation*, or *MMA* — the appropriately annotated orthographic string from which the pronunciation can be deterministically computed. In this case, something like *kostr′a* would be appropriate. Call this *lexical to MMA* transducer L_{word}; such a transducer can be constructed by composing the lexical acceptor D with M so that $L_{word} = D \circ M$. A transducer that maps from the MMA to the standard spelling *kostra* (кострa) would, among other things, simply delete the stress marks: call this transducer S. The composition $L_{word} \circ S$ then computes the mapping from the lexical level to the surface orthographic level, and its inverse $(L_{word} \circ S)^{-1} = S^{-1} \circ L_{word}^{-1}$ computes the mapping from the surface to all possible lexical representations for the text word. A set of pronunciation rules compiled into a transducer (cf. (Kaplan and Kay 1994) and below) P, maps from the MMA to the (surface) phonetic representation; note that by starting with the MMA, rather than with the more abstract lexical representation, the pronunciation rules do not need to duplicate information that is contained in L_{word} anyway. Mapping from a single orthographic word to its

[5] Here, Ë represents a phoneme which surfaces as /jo/ in other forms (e.g. nominative singular). Conversion between a 'flattened' representation of this kind and a hierarchical representation more in line with standard linguistic models of morphology and phonology is straightforward and we will not dwell on this issue here.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

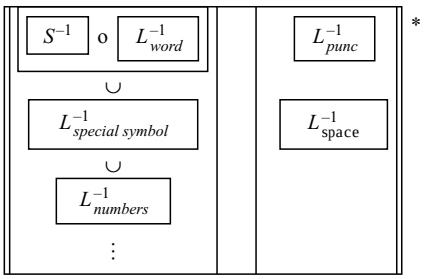


Fig. 1. Overall structure of the lexical analysis portion. Note that SPACE corresponds to whitespace in German or Russian, but ϵ in Chinese or Japanese.

pronunciation thus involves composing the acceptor representing the word with the transducer $S^{-1} \circ L^{-1}_{word} \circ L_{word} \circ P$ (or more fully as $S^{-1} \circ M^{-1} \circ D \circ M \circ P$)[6].

For textual elements such as numbers, abbreviations, and special symbols such as '%', the model just presented seems less persuasive, because there is no aspect of a string, such as '25%' that indicates its pronunciation: such strings are purely *logographic* — some might even argue *ideographic* — representing nothing about the phonology of the words involved[7]. For these cases we presume a direct mapping between all possible forms of *procent*, and the symbol '%': call this transducer L_{perc}, a subset of $L_{special\ symbol}$. Then L^{-1}_{perc} maps from the symbol '%' to the various forms of *procent*. In the same way, the transducer $L^{-1}_{numbers} \cdot L^{-1}_{perc}$ maps from numbers followed by the sign for percent, into various possible (and some impossible) lexical renditions of that string — the various forms to be disambiguated using contextual information, as we shall show later on. Abbreviations are handled in a similar manner: note that abbreviations such as *kg* (кг) in Russian show the same complexity of behavior as *procent*.

So far we have been discussing the mapping of single text words into their lexical renditions. The construction of an analyzer to handle a whole text is based on the observation that a text is simply constructed out of zero or more instances of a text word coming from one of the models described above, i.e. either an ordinary word, an abbreviation, a number, a special symbol, or possibly some combination of numbers with a special symbol; with each of these tokens separated by some combination of whitespace or punctuation. The structure of this model of lexical analysis is summarized in Figure 1.

We presume two models for space and punctuation. The model L^{-1}_{SPACE}, maps between interword SPACE and its potential lexical realizations, usually a word boundary,

[6] It should be stressed that nothing in the Bell Labs TTS system *requires* one to strictly abide by this architecture. One could for example take a morphological FST prepared for some other purpose that converts written text directly to a morphological analysis; and then prepare a separate phonological FST that goes from the morphological analysis to a phonetic transcription. The main advantage of the architecture described here is that the lexical to MMA transduction can in general be *shared* between the surface-to-lexical and lexical to pronunciation transductions, resulting in a reduction of development work.

[7] This is certainly not completely true in all such cases, as 'mixed' representations such as *1st* and *2nd* suggest. But such cases are most easily treated as also being logographic, at least in the present architecture.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

but in some cases a higher-level prosodic phrase boundary. Interword SPACE is parameterized so that in European languages, for example, it corresponds to actual whitespace, whereas in Chinese or Japanese, it corresponds to ϵ. Similarly, the model L_{punc}^{-1} maps between punctuation marks (possibly with flanking whitespace) and the lexical realization of those marks: in many, though not all, cases the punctuation mark may correspond to a prosodic phrase boundary.

The output of the lexical analysis WFST diagramed in Figure 1 is a lattice of all possible lexical analyses of all words in the input sentence. In general we want to remove contextually inappopriate analyses, and to pick the 'best' analysis in cases where one cannot make a categorical decision. This is accomplished by a set of one or more *language model* transducers — henceforth Λ — which are derived from rules and other linguistic descriptions that apply to contexts wider than the lexical word. Phrasal accentuation and prosodic phrasing are also handled by the language model transducers[8]. The output of composing the lexical analysis WFST with Λ is a lattice of contextually disambiguated lexical analyses. The best path of this lattice is then selected, using a shortest path algorithm. Costs on the lattice may be costs hand-selected to disfavor certain lexical analyses — see the Russian percentage example detailed in a subsequent section; or they may be genuine data-derived cost estimates, as in the case of the Chinese lexical analysis WFST, where the costs correspond to the negative log (unigram) probability of a particular lexical entry (Sproat *et al.* 1996). Given the best lexical analysis, one can then proceed to apply the phonological transducer (or set of transducers) P to the lexical analysis, or more properly to the lexical analysis composed with the lexical to MMA map M, as we saw above. Although the lexical to MMA map M was introduced as mapping from the lexical analyses of ordinary words to their MMA, if the map is constructed with sufficient care it can serve as the transducer for lexical analyses coming from any of the text-word models.

3 The tools

The construction of the WFSTs depends upon a lexical toolkit — *lextools* — that allows one to describe linguistic generalizations in linguistically sensible human-readable form. The toolkit has more or less the same descriptive power as the Xerox tools (Karttunen and Beesley 1992; Karttunen 1993), though the current version of lextools lacks some of the debugging capabilities of the Xerox system, and the Xerox tools do not allow costs in the descriptions whereas lextools does[9].

[8] To date our multilingual systems have rather rudimentary lexical class based accentuation rules, and punctuation-based phrasing. Thus these components of the systems are not as sophisticated as the equivalent components of our English system (Wang and Hirschberg 1992; Hirschberg 1993; Sproat 1994). This is largely because the relevant research has not been done for most of the languages in question, rather than for technical problems in fitting the results of that research into the model.

[9] The work is also similar to the Too_{L}iP toolkit for linguistic rules for TTS discussed by van Leeuwen (1989). However, the latter does not compile the rules into (W)FSTs. Instead, the right and left contexts are compiled into an FSM-like format, which is then used to match these contexts for the rules at runtime. Note also that, unlike the current system which compiles linguistic descriptions into WFSTs that allow multiple outputs, Too_{L}iP

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

Some of the tools do not require much comment for readers familiar with previous work on finite state phonology and morphology. In addition to some basic tools to deal with machine labels, there is a tool *compwl* that compiles lists of strings or more general regular expressions into finite state lexicons; a tool *paradigm* to construct inflected morphological forms out of inflectional paradigm descriptions and lexicons that mark the paradigm affiliation of stems; a tool *arclist* for constructing finite state word grammars (the name being inspired by Tzoukermann and Liberman (1990)); and a weighted rewrite rule compiler (Kaplan and Kay 1994), based on the algorithm described in (Mohri and Sproat 1996)[10].

The tool *numbuilder* constructs the transducer $L^{-1}_{numbers}$, which converts strings of digits up to a user-defined length into number-names appropriate for that string. The construction factors the problem of numeral expansion into two subproblems. The first of these involves expanding the numeral sequence into a representation in terms of sums of products of powers of the base — usually ten; call this the *Factorization* transducer. The second maps from this representation into number names using a lexicon that describes the mapping between basic number words and their semantic value in terms of sums of products of powers of the base; call this the *Number Lexicon* transducer. A sample Number Lexicon fragment for German is shown in Figure 2. The first stage of the expansion is language- or at least language-area dependent since languages differ on which powers of the base have separate words (see (Hurford 1975), *inter alia*): so Chinese and Japanese have a separate word for 10^4, whereas most European languages lack such a word. The full numeral expansion transducer is constructed by composing the Factorization transducer with the transitive closure of the Number Lexicon transducer. In some languages, additional manipulations are necessary, and these involve the insertion of a special 'filter' transducer between the Factorization and Number Lexicon transducers. In German, for example, the words for decades come after the words for units: so $34 = 3 \cdot 10^1 + 4$ becomes *vierunddreißig* (four+and+thirty), thus suggesting the order $4 + 3 \cdot 10^1$. This reversal can be accomplished by a filter transducer, which for lack of a better name we will call *Decade Flop*. The construction of the numeral expander for German is shown schematically in Figure 3. (Note that the insertion of *und* in examples like *vierunddreißig* 'thirty four' is not actually accomplished by numbuilder: such 'clean up' operations can be handled using rewrite rules.)

One tool that has so far not been applied in the system is the decision tree compiler, described in Sproat and Riley (1996). We are hoping to apply this in the future for, among other things, phrasing models of the kind discussed in Wang and Hirschberg (1992).

functions in a completely deterministic fashion, in that for any given input there can only be one output.

[10] Weights can be incorporated into rewrite rules by simply annotating the output ('ψ') portion of the rule with weights. So a rule such as $a \rightarrow b < 2.0 > \ / \ ca__d$ states that a in the environment $ca__b$ is rewritten as b with cost 2.0.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

```
/{1}                    :   ('ein{num}({masc}|{neut}){sg}{##})/
/{1}                    :   ('eine{num}{femi}{sg}<1.0>{##})/
/{2}                    :   (zw'e{num}{##})/
                        ⋮
/{0}{+++}{1}{10Λ1})     :   (z'ehn{num}{##})/
/{1}{+++}{1}{10Λ1})     :   ('elf{num}{##})/
/{2}{+++}{1}{10Λ1})     :   (zw'ölf{num}{##})/
                        ⋮
/{2}{10Λ1})             :   (zw'an{++}zig{num}{##})/
/{3}{10Λ1})             :   (dr'ei{++}ßig{num}{##})/
                        ⋮
/{10Λ2})                :   (h'under{num}{##})/
/{10Λ3})                :   (t'ausend{num}{neut}{##})/
```

Fig. 2. German number lexicon. The cost of 1.0 on the feminine form *eine* effectively disfavors this form, so that it will only be selected in appropriate predefined contexts.

$$
\begin{array}{llll}
234 & \circ\ \boxed{\text{Factorization}} & \Rightarrow & 2.10^2 + 3.10^1 + 4 \\
 & \circ\ \boxed{\text{Decade Flop}} & \Rightarrow & 2.10^2 + 4 + 3.10^1 \\
 & \circ\ \boxed{\text{Number Lexicon}}^{*} & \Downarrow & \\
 & \textit{zwei+hundert+vier+und+dreißig} & &
\end{array}
$$

Fig. 3. Expansion of *234* in German using *numbuilder*.

4 Russian percentages

Let us return to the example of Russian percentage terms. Assume that we start with a fragment of text such as с 5% скидкой *s 5% skidkoj* (with 5% discount) 'with a five-percent discount'. This is first composed with the lexical analysis WFST to produce a set of possible lexical forms; see Figure 4. By default the lexical analyzer marks the adjectival readings of '%' with '$\star$', meaning that they will be filtered out by the language-model WFSTs, if contextual information does not save them. Costs on analyses (here represented as subscripted floating-point numbers) mark constructions — usually oblique case forms — that are not in principle ill-formed but are disfavored except in certain well-defined contexts. The correct analysis (boxed in Figure 4), for example, has a cost of 2.0 which is an arbitrary cost assigned to the oblique instrumental adjectival case form: the preferred form of the adjectival rendition '%' is masculine, nominative, singular if no constraints apply to rule it out.

Next the language model WFSTs Λ are composed with the lexical analysis lattice. The WFSTs Λ include transducers compiled from rewrite rules that ensure that the adjectival rendition of '%' is selected whenever there is a noun following the percent expression, and rules that ensure the correct case, number and gender of the adjectival form given the form of the following noun. In addition, a filter expressable as $\neg(\Sigma^* \boxed{\star} \Sigma^*)$ removes any analyses containing the tag '$\star$'. See Figure 5. The best-cost analysis among the remaining analyses is then selected. Finally, the lexical analysis is composed with $M \circ P$ to produce the phonetic transcription: `s # PiT"!p~r@c"Entn&y # sK"!tk&y`.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

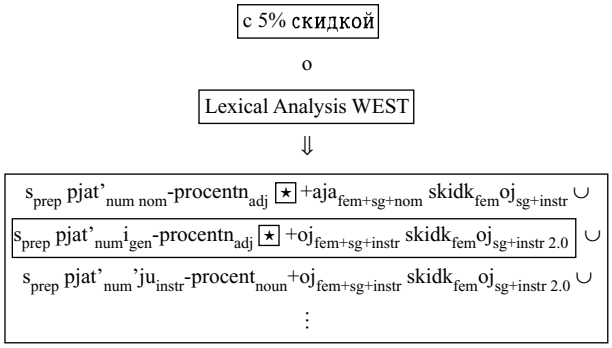


Fig. 4. Composition of с 5% скидкой *s 5% skidkoj* 'with a 5% discount' with the lexical analysis WFST. By default the adjectival readings of '%' are marked with '⋆', which means that they will be filtered out by the language-model WFSTs; see Figure 5. The boxed analysis is the correct one. Costs on analyses mark constructions – usually oblique case forms – that are not in principle ill-formed but are disfavored except in certain well-defined contexts.

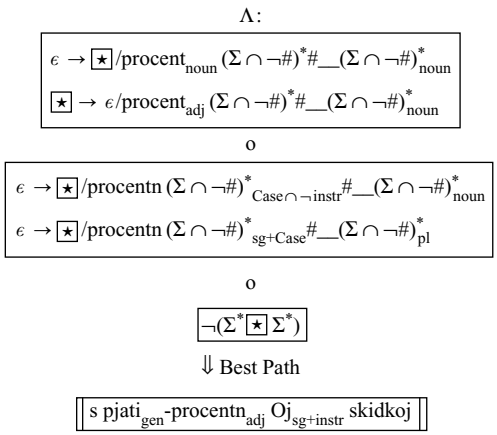


Fig. 5. A subset of the language model WFSTs related to the rendition of percentages. The first block of rules ensures that adjectival forms are used before nouns, by switching the tag '⋆' on the adjectival and nominal forms. The second block of rules deals with adjectival agreement with the adjectival forms. The final block is a filter ruling out the forms tagged with '⋆'. The (correct) output of this sequence of transductions is shown at the bottom.

5 Size and speed issues

Table 1 gives the sizes of the lexical analysis WFSTs for the languages Spanish, German, Russian, Mandarin and Japanese. To a large extent, these sizes accord with our intuitions of the difficulties of lexical processing in the various languages. So Russian is very large, correlating with the complexity of the morphology in that language. German is somewhat smaller. Mandarin has a small number of states, correlating with the fact that Mandarin words tend to be simple in terms of morphemic structure; but there are a relatively large number of arcs, due to the large character set involved. Sizes for the Spanish transducer are misleading since

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

Table 1. *Sizes of lexical analysis WFSTs for selected languages*

	States	Arcs
Spanish	8,602	17,236
German	74,185	310,233
Russian	139,592	495,847
Mandarin	48,015	278,905
Japanese	152,176	883,650

the current Spanish system includes only minimal morphological analysis: note, though, that morphological analysis is mostly unnecessary in Spanish for correct word pronunciation.

While the transducers can be large, the performance (on an SGI Indy or Indigo) is acceptably fast for a TTS application. Slower performance is observed, however, when the system is required to explore certain areas of the network, as for example in the case of expanding and disambiguating Russian number expressions.

To date, no formal evaluations have been performed on the correctness of word pronunciation in the various languages we are working on, largely because there is still work to be done before the systems can be called complete. An evaluation of the correctness of word segmentation in the Mandarin Chinese system is reported in (Sproat *et al.* 1996).

6 Summary and future work

The system for text analysis presented in this paper is a complete working system that has been used in the development of working text analysis systems for several languages. In addition to German, Spanish, Russian, Mandarin and Japanese, systems for Romanian, Italian and French have been built. Needless to say, the development time necessary to produce a system with an acceptable level of performance depends greatly upon the language in question and the number of lexical resources available to begin with. For instance, a text analysis module was built for Italian in a few person-weeks, making use of large online dictionary for Italian which provided pronunciations for inflected forms. In contrast, Russian has required a couple of years of development, despite the fact that we had electronic dictionaries of the language.

From the point of view of previous research on linguistic applications of finite state transducers, some aspects of this work are familiar, some less so. Familiar, of course, are applications to morphology, phonology, and syntax, though most previous work in these areas has not made use of *weighted* automata. More novel are the applications to text 'preprocessing', in particular numeral expansion and word segmentation. From the point of view of text analysis models for text-to-speech, the approach is quite novel since, as described in the introduction, most previous work treats certain operations, such as word segmentation or numeral expansion in a

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

preprocessing phase that is logically prior to the linguistic analysis phase; we have argued here against this view.

Two areas of future work both depend upon an important property of the FSM toolkit on top of which the lextools toolkit is built. Underlying the notion of an FSM is the more general notion of a *generalized* state machine (GSM). An important property of GSMs is that it is not necessary to know beforehand which arcs leave a given state; rather one can construct just the arcs one needs 'on the fly' as one is using the machine, for example in a composition with another machine. This has two important consequences. First of all, for a strictly *finite* state machine, it is not necessary to explicitly construct the machine beforehand, and this in turn implies that one can avoid precompiling very large FSMs, so long as one can provide an algorithm for constructing the machine on the fly. One example is in discourse analysis, where one wants to remember which words or lemmata one has already seen; as previous work on accenting suggests (Hirschberg 1993), this kind of information is useful for TTS, and is in fact used in the American English version of the Bell Labs synthesizer. In theory, assuming the set of words or morphological stems is closed, one could construct an FSM that would 'remember' when it had seen a word; needless to say, such a machine would be astronomical in size, and so the precompilation of this machine is out of the question; one could however envision dynamically building states that 'remember' that a particular word has been seen. Secondly, one can in principle construct GSMs which have greater than finite state power, again providing that one can specify an algorithm for constructing on the fly the arcs leaving a given state. One obvious example is a 'copy machine' which will recognize which strings from a lattice have the property that they are of the form ww, for some string w; this problem comes up in the analysis of morphological reduplication. Precompiling such a machine as an FSM for copies of unbounded length is of course impossible; however, it is possible to construct a GSM which can be composed with an arbitrary (acyclic) lattice and will find exactly those strings with the desired property. Future work on the text analysis model presented here will focus in part on the application of generalized state machines to various linguistic problems.

Acknowledgments

I wish to acknowledge Michael Riley, Fernando Pereira and Mehryar Mohri of AT&T Laboratories, without whose underlying FSM toolkit the lexical toolkit and text analysis work reported here would not have been possible. Kimmo Koskenniemi provided some valuable comments on an earlier version of this paper. Bernd Möbius implemented the German text analysis example cited in the text.

References

Daniels, Peter and Bright, William. 1996. *The World's Writing Systems*. Oxford University Press, New York.

Franks, Steven. 1994. Parametric properties of numeral phrases in Slavic. *Natural Language and Linguistic Theory*, **12**: 597–674.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press

Hirschberg, Julia. 1993. Pitch accent in context: Predicting intonational prominence from text. *Artificial Intelligence*, **63**: 305–340.

Hurford, James. 1975. *The Linguistic Theory of Numerals*. Cambridge University Press, Cambridge.

Kaplan, Ronald and Kay, Martin. 1994. Regular models of phonological rule systems. *Computational Linguistics*, **20**: 331–378.

Karttunen, Lauri and Beesley, Kenneth. 1992. Two-level rule compiler. Technical Report P92–00149, Xerox Palo Alto Research Center.

Karttunen, Lauri, Kaplan, Ronald, and Zaenen, Annie. 1992. Two-level morphology with composition. In *COLING-92*, pp. 141–148. COLING.

Karttunen, Lauri, Chanod, Jean-Pierre, Grefenstette, Gregory, and Schiller, Anne. 1997. Regular expressions for language engineering. *Natural Language Engineering*, this issue.

Karttunen, Lauri. 1993. Finite-state lexicon compiler. Technical Report P93–00077, Xerox Palo Alto Research Center.

Koskenniemi, Kimmo. 1983. *Two-Level Morphology: a General Computational Model for Word-Form Recognition and Production*. Ph.D. thesis, University of Helsinki, Helsinki.

Mohri, Mehryar and Sproat, Richard. 1996. An efficient compiler for weighted rewrite rules. In *34th Annual Meeting of the Association for Computational Linguistics*, pp. 231–238, Morristown, NJ. Association for Computational Linguistics.

Pereira, Fernando and Riley, Michael. 1996. Speech recognition by composition of weighted finite automata. CMP-LG archive paper 9603001. `xxx.lanl.gov/ps/cmp-lg/9603001`.

Pereira, Fernando, Riley, Michael, and Sproat, Richard. 1994. Weighted rational transductions and their application to human language processing. In *ARPA Workshop on Human Language Technology*, pp. 249–254. Advanced Research Projects Agency, March 8–11.

Sproat, Richard and Riley, Michael. 1996. Compilation of weighted finite-state transducers from decision trees. In *34th Annual Meeting of the Association for Computational Linguistics*, pp. 215–222, Morristown, NJ. Association for Computational Linguistics.

Sproat, Richard, Shih, Chilin, Gale, William, and Chang, Nancy. 1996. A stochastic finite-state word-segmentation algorithm for Chinese. *Computational Linguistics*, **22**

Sproat, Richard. 1992. *Morphology and Computation*. MIT Press, Cambridge, MA.

Sproat, Richard. 1994. English noun-phrase accent prediction for text-to-speech. *Computer Speech and Language*, **8**: 79–94.

Traber, Christof. 1995. SVOX: The implementation of a text-to-speech system for German. Technical Report 7, Swiss Federal Institute of Technology, Zurich.

Tzoukermann, Evelyne and Liberman, Mark. 1990. A finite-state morphological processor for Spanish. In *COLING-90, Volume 3*, pp. 277–286. COLING.

van Leeuwen, Hugo. 1989. *Too$_L$iP: A Development Tool for Linguistic Rules*. Ph.D. thesis, Technical University Eindhoven.

Voutilainen, Atro. 1994. Designing a parsing grammar. Technical Report 22, University of Helsinki.

Wang, Michelle and Hirschberg, Julia. 1992. Automatic classification of intonational phrase boundaries. *Computer Speech and Language*, **6**: 175–196.

https://doi.org/10.1017/S1351324997001654 Published online by Cambridge University Press