

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47

8

A Survey of Discriminative Language Modeling Approaches for Large Vocabulary Continuous Speech Recognition

Brian Roark

Discriminative training of language models has been the focus of renewed research interest recently, and significant improvements to high accuracy large vocabulary continuous speech recognition have been demonstrated through the use of such models. In this chapter, we review past and present work on this topic, and focus in particular on three key issues: training data, learning algorithms and features. We will show how simple models trained under standard approaches can provide significant accuracy gains with negligible impact on decoding time, and how richer models can provide further accuracy gains. We will argue that while the published results are promising, many practical issues remain under-explored, so that further improvements might be expected as the best practices become more developed.

8.1 Introduction

It has long been the case that system improvements in statistical speech recognition have come far more from acoustic modeling innovations than from language modeling innovations – to the extent that some top-notch researchers have quit working on language modeling for speech recognition, sometimes with a parting message of warning to future researchers against expecting system improvements. (See the extended version of (Goodman 2001),

01 which can be found with the search engine query: ‘language modeling’ ‘all hope abandon.’)
02 There are a number of reasons why little progress was observed due to language modeling,
03 including (though certainly not confined to): extreme sparse data; optimization of objectives
04 only very loosely corresponding to the true system objective; and difficulty integrating
05 disparate potential features into a single model. Sparse data is a problem that can be alleviated
06 with increased training data, which some may not consider a particularly interesting solution,
07 though it can require solving difficult problems that arise when scaling up algorithms to
08 handle larger data sets. The continued improvement of performance in speech recognition
09 and machine translation due to language model training on extremely large training sets
10 (Emami *et al.* 2007; Popat *et al.* 2007) illustrates the extent to which sparse data has been
11 impacting language models. Even what have been considered very large corpora in the past do
12 not provide adequate training for multinomial n -gram models with millions or even billions
13 of parameters. In Popat *et al.* (2007), for example, the models have hundreds of billions
14 of parameters (corresponding to n -gram features of various orders), which are trained on
15 corpora of several trillion tokens; they find a linear gain with each doubling of the training
16 data, with no indication of a plateau.

17 In this chapter, we will focus on language modeling improvements in large vocabulary
18 continuous speech recognition (LVCSR) that are not due to an increase in the amount of
19 training data, but rather to improvements in the other two dimensions of the problem that were
20 identified in the previous paragraph: optimization of parameters with respect to objectives
21 more closely related to actual system objectives; and the inclusion of a large number of
22 heterogeneous feature sets derived from the output string hypotheses.

23 Neither of these two considerations is new to research in language modeling for LVCSR.
24 It was suggested in Jelinek (1996) to train an ‘acoustic-sensitive’ language model whose
25 parameters would be trained to minimize the expected uncertainty of the spoken text given the
26 acoustic signal. Discriminative training methods that were successfully pursued for acoustic
27 modeling, such as MMIE (Bahl *et al.* 1986) or related error-corrective training methods
28 (Bahl *et al.* 1993; Juang *et al.* 1997), spurred initial research into discriminative modeling
29 approaches such as Chen *et al.* (2000); Stolcke and Weintraub (1998); Stolcke *et al.* (2000),
30 with modest utility. In addition, generative maximum entropy models have been explored for
31 combining diverse features within a single model (Khudanpur and Wu 2000; Rosenfeld *et al.*
32 2001), including n -grams, topic labels and syntactic features. Thus the notion of exploiting
33 diverse feature sets is not new to the field; and neither are discriminative approaches to
34 parameter estimation.

35 Over the past few years, however, there has been a renewed effort in this direction, driven
36 in large part by developments in the machine learning community. The benefits of global
37 conditional modeling for sequence learning (Lafferty *et al.* 2001) and the application of large
38 margin methods to sequence processing through the Perceptron algorithm (Collins 2002;
39 Freund and Schapire 1999) have provided a unifying framework within which many related
40 methods can be pursued and compared. The learning and inference algorithms associated
41 with some of these approaches are efficient enough to be scaled up, so that models making
42 use of very large feature sets can be effectively trained and used for inference. Well-motivated
43 and theoretically grounded regularization techniques have been shown to effectively control
44 for the overtraining that language models have long suffered from. Strong empirical gains
45 have been realized on difficult tasks and continue to be realized up to the present within this
46 paradigm.

47

01 This chapter addresses discriminative language modeling approaches in general, not
 02 kernel approaches in particular. The approaches that are the focus of this chapter do not
 03 represent the only strand of research focused on leveraging competing LVCSR system
 04 hypotheses to achieve reductions in error rate. The approaches we describe here are focused
 05 on *training* models with a discriminative objective function, not *inference* techniques for
 06 optimizing an objective. Hence approaches such as Minimum Bayes-Risk classification (Goel
 07 and Byrne 2000) or confusion networks (‘sausages’) (Mangu *et al.* 2000) will not be covered
 08 here, although methods exist for learning their edit distance (Shafran and Byrne 2004)¹. Note
 09 that inference using such techniques can also be pursued with models trained as presented
 10 here, hence the techniques are complementary. We will also omit any detailed discussion
 11 of methods that perform post-processing on LVCSR output to correct specific errors, such
 12 as induction of transformational rules on the output (Mangu and Padmanabhan 2001) or
 13 methods for proposing alternatives that are not present in the system output and choosing
 14 between them (Ringger and Allen 1996; Zhou *et al.* 2006b). Again, such methods could
 15 also be applied to the output of systems using models trained with methods presented in this
 16 survey, hence they are complementary to what will be discussed here. Also, we have chosen to
 17 focus on large vocabulary systems, hence will not discuss the application of similar methods
 18 in small vocabulary applications, e.g. that in Kuo *et al.* (2002).

21 8.2 General Framework

22
 23 Most of the work covered in this chapter falls within the general framework of linear models,
 24 which we present here largely following the notation of Collins (2002), much as presented
 25 in (Roark *et al.* 2007). These approaches assume the presence of several components, as
 26 follows.

27 First, these are supervised approaches, hence requiring the presence of training data,
 28 consisting of a set of m input sequences $\bar{x}_i \in \mathcal{X}$ and corresponding reference (ground truth)
 29 output sequences $\bar{y}_i^{\text{ref}} \in \mathcal{Y}$, where \mathcal{X} denotes the set of possible inputs and \mathcal{Y} the set of
 30 possible outputs. In the case of language modeling for LVCSR, using a vocabulary Σ , the
 31 input set \mathcal{X} consists of acoustic samples, the output set consists of all possible transcriptions,
 32 i.e. $\mathcal{Y} = \Sigma^*$, and the reference (ground truth) output sequence for any given utterance is the
 33 manual (human) transcription of that utterance.

34 Next, for each input sequence $\bar{x} \in \mathcal{X}$, the approaches require some mechanism for
 35 identifying a set of output candidates, which is often denoted as a function $\mathbf{GEN}(\bar{x}) \subseteq \mathcal{Y}$.
 36 For example, $\mathbf{GEN}(\bar{x})$ could be the word lattice output or n -best list from a baseline LVCSR
 37 system.

38 Third, there must be a mapping from each input:output pairing $(\bar{x}, \bar{y}) \in \mathcal{X} \times \mathcal{Y}$ to a real
 39 valued d -dimensional feature vector $\psi(\bar{x}, \bar{y}) \in \mathbb{R}^d$, for some d . In the current case, this could
 40 be a mapping from a string to the set of n -grams of order $\leq k$ that occur in the string. An
 41 example feature would then be something like ‘feature 1236: number of times the bigram
 42 “dog house” occurs in the string’. Each explicit n -gram to be included in the model would be
 43 a separate feature in the vector.

44
 45 ¹The methods for learning the edit distance used in such techniques are focused on learning a weighted
 46 transduction between reference transcripts and LVCSR system outputs. This learned transduction can be seen as
 47 a form of discriminative language modeling, since it effectively moves the centroid closer to the reference.

01 Finally, there must be a real valued parameter vector $\alpha \in \mathbb{R}^d$ of the same dimension d as
 02 the feature vector. For example, feature 1236 mentioned above may have a parameter value
 03 of 0.2 assigned to it. Methods for determining the parameter values associated with each
 04 feature are perhaps the key difference between many of the different approaches that will be
 05 presented here, though they are all discriminative rather than generative approaches.

06 Given these four components, we can define the best scoring (or 1-best) output \bar{y}_i^{\max} for a
 07 given input x_i as

$$08 \quad \bar{y}_i^{\max} = \arg \max_{\bar{y} \in \text{GEN}(\bar{x}_i)} \psi(\bar{x}_i, \bar{y}) \cdot \alpha, \quad (8.1)$$

09 where $\psi(\bar{x}_i, \bar{y}) \cdot \alpha$ is the dot product.

10 In the remainder of this section, we will focus on each of these four components in turn,
 11 highlighting some of the key trends and variations in the work from this area.

12 8.2.1 Training Data and the GEN Function

13 Training data for discriminative language modeling consists of input speech recordings and
 14 reference output strings representing the gold standard transcription of the speech. The role
 15 of the GEN function is typically taken by a baseline LVCSR system, so that $\text{GEN}(\bar{x})$ for
 16 some input speech is an n -best list or word lattice. Hence most discriminative language
 17 modeling approaches fall within the range of approaches more traditionally known as
 18 n -best or lattice rescoring. Working from lists and lattices, however, is not a requisite part of
 19 the approach, and some recent papers (Kuo *et al.* 2007; Lin and Yvon 2005) have modified
 20 the weights directly within the baseline recognition system, rather than adopting a rescoring
 21 or reranking paradigm. See section 8.2.2 for further details on directly updating the baseline
 22 recognition system, where it is discussed as a method for encoding features to allow for
 23 efficient processing.

24 The output of LVCSR systems can be thought of as weighted, acyclic finite state automata,
 25 with words and weights labeling the arcs. Figure 8.1 shows three kinds of representation of
 26 weighted output from an LVCSR system. For all such weighted automata, the semiring must
 27 be specified, so that the weights can be interpreted. The automata in Figure 8.1 are in the
 28 Real semiring, meaning that the weights along a path are multiplied, and when two paths
 29 are combined, their weights are added, which is the appropriate semiring when dealing with
 30 probabilities². In the figure, the weights define a conditional probability distribution over
 31 three possible strings: ‘a b c’ with probability 0.3; ‘a g c’, with probability 0.3; and ‘e d c
 32 c’, with probability 0.4. Although the weights are presented as conditional probabilities, they
 33 can be arbitrary.

34 Not all weighted finite-state automata (WFA) can be determinized, but any acyclic WFA
 35 can be determinized (Mohri and Riley 1997). The result of determinization on the n -best
 36 list automaton in Figure 8.1(a) is the deterministic tree in Figure 8.1(b). As with the original
 37 representation, a candidate string is represented by the concatenation of symbols labeling
 38 the arcs of any path from the start state at the left to a final state denoted with a double

39 ²The most common semirings in speech processing are the Tropical and Log semirings, both of which are
 40 appropriate for negative log probabilities. Both add weights along the path. The Tropical semiring takes the **min** of
 41 two weights when combining paths, and the Log semiring takes the log of the sum of the exponentials: $A \oplus B =$
 42 $-\log(\exp(-A) + \exp(-B))$.




Figure 8.1 Three weighted, acyclic finite-state automata representations of the output from a speech recognizer: (a) n -best list; (b) deterministic tree; (c) deterministic word lattice. The automata are in the Real semiring, and all three represent the same weighted distributions over the three strings.

circle. What is key to notice from these figures is that determinization preserves the path weight of every path in the automaton. Recall that weights are combined along a path in the Real semiring by multiplying. Inspection will verify that all three of the paths in the original automaton are present in the deterministic tree representation with the same path weight.

Deterministic WFA can be minimized, i.e. converted to a representation of exactly the same weighted paths, but with a minimal number of states and arcs to represent it (Mohri and Riley 1997). For the current example, this minimized WFA is shown in Figure 8.1(c). These are word lattices, which are simply a more compact representation of a set of distinct weighted hypotheses output from the LVCSR system. Efficient use of this kind of compact representation allows for effective rescoring in scenarios where brute force enumeration of all distinct candidates would be intractable. Effective use of these representations, however, precludes the use of global features, unless they can be decomposed for dynamic programming. Some approaches to discriminative language modeling, e.g. those making use of long-distance dependencies uncovered via context-free parsing, will be forced to operate on undeterminized or unminimized n -best representations, which limits the number of candidates that can be tractably scored.

There are a couple of reasons to introduce WFA representations in this context. First, they become important in the next section, when efficient feature mapping is discussed. Also, it should be clear from the discussion that there is no difference between n -best lists and word lattices, other than determinization and minimization – they both represent sets of competing (weighted) candidate transcriptions. Even 1-best output of a system is a special case of these acyclic WFA. Hence we can discuss the approaches in general, without having to specify whether the data is in the form of n -best lists or lattices. We will thus henceforth refer to them simply as lattices, which may or may not be determinized or minimized. General operations on WFA can be used to specify operations on the lattices, including intersection, composition with a transducer, and finding the minimum cost (or highest scoring) path.

Note that working within a reranking or rescoring paradigm has a few immediate ramifications for the general modeling framework we are presenting. Most critically, it is not guaranteed in this approach that the reference transcription is present in the lattice. Second, if training lattices are produced for utterances that were also used for building the models of the baseline recognizer, this may introduce a mismatch between training conditions and testing conditions. These are two key issues investigated in Roark *et al.* (2004, 2007), and we discuss them here in turn.

(a) Reference Transcript

If the reference transcript for a particular utterance is not in the lattice output in the training data, then standard parameter estimation algorithms will be trying to achieve an objective that cannot be reached. These parameter estimation algorithms will be discussed in section 8.2.3, but, intuitively, an algorithm (such as the Perceptron algorithm) that penalizes candidates other than the gold standard (reference) will move parameters to penalize every candidate if the gold standard is not in the set of candidates, even the highest accuracy candidate in the whole set.

One quick fix to this problem would be to append the reference transcript to the set of candidates in the training set. The main problem with this is that it creates a great mismatch between the training condition and the testing condition, in that the training examples are being given better candidate sets than will be seen under testing conditions. Of lesser significance is the problem of acquiring features derived from the baseline LVCSR system for the reference transcript, if it was not one of its candidates. For example, the baseline system score is typically a useful feature in the model. Obtaining the baseline system score for the reference may require having some access to the baseline system models, to constrain the decoder to output the reference.

A better solution is to define the gold standard candidate in terms of maximum accuracy or minimum error rate. Then, if the reference is in the candidate set, it will be the gold standard; otherwise, another candidate from the set is chosen as the gold standard, to avoid the problems with parameter estimation. When the reference is present in the lattice, the gold standard is unique. If the reference is not present in the lattice, then there may not be a unique gold standard: more than one candidate may have the minimum error rate. The results in Roark *et al.* (2004, 2007) empirically validate the utility of using a minimum error rate gold standard versus a reference gold standard in the Perceptron algorithm. In those papers, one gold standard is chosen from among the set of minimum error rate candidates; in (Singh-Miller and Collins 2007), all members from the set were collectively used as the gold standard in a loss-sensitive version of the Perceptron algorithm. See section 8.2.3 for more details about parameter estimation when the reference is not in the lattice.

Note that this is also an issue when training outside of a rescoring or reranking paradigm, as noted in Kuo *et al.* (2007), even with known training data. If the vocabulary is pruned, there will be out-of-vocabulary (OOV) terms in the references which the closed vocabulary LVCSR system cannot recognize. One method for decreasing the OOV rate in any of these training scenarios, however, is to map terms to ‘equivalent’ items that are in the vocabulary, e.g. legal spelling variations.

(b) Seen Versus Unseen Data

The second issue that was mentioned regarding the training data is the potential mismatch between training and testing scenarios when training lattices are produced on previously seen data. This is an issue with any kind of a reranking problem, and similar solutions to the one that will be discussed here have been used for other such problems, e.g. (Collins 2000). In fact, the problem is to minimize mismatch, because most problems have a constrained amount of supervised training data (with both input and output sides of the mapping), hence the same data will often be used both for training the baseline models and for training discriminative language models. It is well-known that n -gram language models are overparameterized and

01 overtrained; the result of this is that the LVCSR output for seen data will have fewer errors
02 than the output for unseen data, and those errors that do occur may differ from the kinds of
03 error observed for unseen data. The discriminative reranking model will then learn to solve a
04 different problem than the one that will be encountered at test time. To obtain a useful model,
05 the mismatch needs to be reduced by producing training lattices in an unseen condition.

06 An unseen condition can be achieved through cross-validation, whereby the training data
07 is split into k subsets, and the baseline models for each subset j are trained on the other
08 $k-1$ subsets $(1, \dots, j-1, j+1, \dots, k)$. Note that this can still result in a mismatch,
09 because the testing condition will use baseline models trained on all k subsets. The hope
10 is that, if k is large enough, the difference in performance when training on the extra section
11 will be relatively small, particularly compared with the difference in seen versus unseen
12 performance. Again, the results in Roark *et al.* (2004, 2007) empirically validate the utility of
13 cross-validating the baseline language model training: not cross-validating resulted in models
14 that achieved no gain over the baseline system, whereas cross-validation produced significant
15 word-error rate (WER) reductions of over 1% absolute.

16 While the presentation here has been on the importance of cross-validating the models
17 used in the baseline systems, previous work has only examined the utility of cross-validating
18 the language models, not the acoustic models. The reasons for this are twofold. First,
19 conventional wisdom is that acoustic models, by virtue of having far fewer parameters than
20 language models, tend not to overtrain to the same degree, so that the seen/unseen difference
21 in performance is relatively small. Second, while baseline n -gram models are relatively
22 easy and fast to train (even for large corpora), since they typically use relative frequency
23 estimation and simple smoothing to estimate the parameters, acoustic model training is a far
24 more intensive process, so that training a baseline acoustic model on a large training corpus
25 may take many days or more on multiple processors. Hence the cost of cross-validation is
26 high and the perceived benefit low. Even so, it would be of interest to know exactly how
27 much mismatch exists between seen and unseen data with respect to the acoustic model,
28 and whether this may actually impact the utility of training data for discriminative language
29 modeling.

30

31

32 8.2.2 Feature Mapping

33

34 The third component of the general framework that we are presenting is the mapping
35 from input/output pairs (\bar{x}, \bar{y}) to a feature vector $\psi(\bar{x}, \bar{y}) \in \mathbb{R}^d$. In the current context,
36 there are two scenarios that we have discussed: reranking from lattices; and updating the
37 baseline LVCSR system directly. We will first discuss feature mapping for reranking, where
38 candidates are represented as an acyclic WFA of candidates; then for direct update of the
39 system.

40

41 For reranking, one approach to feature mapping is to simply perform feature mapping for
42 each candidate transcript independently. For example, if an existing statistical parser were to
43 be applied to the strings and the result used to extract features, then the particular parser may
44 not be able to handle anything other than a string at a time. This requires explicit enumeration
45 of all candidates, which will limit the number of candidates that can be considered for any
46 given utterance.

47

48 A more efficient strategy is to take advantage of shared structure in a word lattice to extract
49 features without having to explicitly enumerate every path in the lattice. This is particularly

01 important in LVCSR, where the number of competing hypotheses can be in the hundreds
 02 of thousands or millions, and the feature set in the several millions. N -gram features are
 03 popular in large part because they have an efficient finite-state structure, and this structure
 04 can be exploited for efficient feature mapping that allows for large lattices to be rescored,
 05 rather than small n -best lists. We will use a toy example to illustrate this. For more details on
 06 n -gram model structure, see Allauzen *et al.* (2003) and Roark *et al.* (2004, 2007, 2004).

07 For simplicity, assume a vocabulary $\Sigma = \{a,b,c\}$, and all strings consist of one or more
 08 symbols from that vocabulary in sequence (Σ^+). The output symbols will be feature indices,
 09 which we will denote f_k for feature with index k . For this example, let us assume the following
 10 unigram and bigram features: $\{f_1 = a; f_2 = b; f_3 = \langle s \rangle a; f_4 = \langle s \rangle b; f_5 = aa; f_6 = ab; f_7 =$
 11 $a \langle /s \rangle; f_8 = ba; f_9 = ba; f_{10} = b \langle /s \rangle\}$. Note that the begin-of-string symbol ($\langle s \rangle$) and end-
 12 of-string symbol ($\langle /s \rangle$) are implicit in the string and not part of the vocabulary. One can see
 13 by inspection of this set that there are all unigram and bigram features including the symbols
 14 ‘a’ and ‘b’, but none including the symbol ‘c’, hence those n -grams will not map to any
 15 feature in the set.

16 The finite-state transducer in Figure 8.2(a) maps from strings in $\{a,b,c\}^+$ to feature indices
 17 from this set. The start state (denoted with \triangleright to its left) implicitly encodes the begin-of-string
 18 symbol, hence the first symbol in the string will be in a bigram beginning with $\langle s \rangle$. There
 19 are two special symbols in the transducer: ϕ , which represents a failure transition; and ϵ ,
 20 which represents the empty string. An arc labeled with ϕ is traversed if and only if there is
 21 no arc leaving the state that is labeled with the next symbol in the input. No input symbols
 22 are consumed or output symbols emitted when the ϕ transition is traversed. An arc with
 23 an ϵ symbol labeling the input side can always be traversed without consuming an input
 24 symbol, while outputting the symbol on the output side. An arc with an ϵ symbol labeling the
 25 output side produces no output while consuming the symbol on the input side. When final
 26 states (denoted with double circle) are reached and there is no more input, the transduction is
 27 successful.

28 Let us step through a couple of examples. First the simple string ‘acb’, which we will
 29 compose with transducer T_1 from Figure 8.2(a). Beginning at the start state, there is an arc
 30 labeled with ‘a’ on the input side, with f_3 on the output side. From the destination state of
 31 that arc, there is no arc with ‘c’ on the input side, so the failure transition (labeled with ϕ)
 32 is traversed. From the destination state of that arc, there is an arc with ‘c’ on the input side,
 33 which is a looping arc with ϵ (empty string, i.e. no feature) on the output side. The transition
 34 from that state with ‘b’ on the input side has f_2 on the output side. The destination state of
 35 that arc is not a final state, but there is a transition with an ϵ on the input side which can be
 36 traversed, which gives f_{10} on the output side and reaches a final state. Hence this transduction
 37 gives the features: f_3 , f_2 and f_{10} . This is not, however, the total feature set, but rather the
 38 highest order n -grams reached at each word. For example, feature f_3 is the bigram ‘ $\langle s \rangle a$ ’,
 39 but there should also be the unigram feature f_1 associated with the ‘a’. In order to derive
 40 the total feature set based on the highest order features, we compose the output sequence
 41 to the transducer T_2 in Figure 8.2(b). In that transducer, all features map to themselves, but
 42 additionally all bigram features ending in an ‘a’ or ‘b’ then also map to the appropriate
 43 unigram feature. The complete transduction from a string S to the feature set is $S \circ T_1 \circ T_2$,
 44 where \circ denotes the composition of finite-state transducers.

45 Note that these are both deterministic on the input side, i.e. every string in $\{a,b,c\}^+$
 46 has one and only one path through $T_1 \circ T_2$. It should be straightforward for the reader
 47




Figure 8.2 String sequence to feature set transducers: (a) from symbols to highest order features; (b) from highest order features to total set of features.

to verify that the input string ‘bbacb’, when composed with $T_1 \circ T_2$, yields the feature sequence $f_4 f_2 f_9 f_2 f_8 f_1 f_2 f_{10}$. Because this feature mapping uses standard finite-state transducer transduction, word lattices encoded as finite-state automata can be composed with these transducers, to produce output sequences for every path in the lattice corresponding to features in the feature set. Multiple paths that share part of their structure will not need to be mapped independently of each other, as they would be in non-deterministic representations of n -best lists.

Note that, at test time, the feature identities are not of interest, but rather the scores of the paths in the lattices, so that the best scoring candidate can be found. There is also a compact finite-state automaton that can score the lattice with n -gram parameter weights, without requiring transduction to the feature labels. The automaton in Figure 8.3 is in the Tropical semiring (sum weights along a path), and makes use of roughly the same topology as the transducer in Figure 8.2(a). Note that all arcs corresponding to bigram features must have the sum of both the bigram feature weight and the corresponding unigram feature weight. Parameter weights associated with bigrams containing the end-of-string symbol are encoded in the final weight of the corresponding final state. Similar structure is used when going to n -grams of higher order than bigrams. See Roark *et al.* (2004, 2007, 2004) for more information about such finite-state n -gram model encoding.

Finite-state transducer encoding of the model is also a central topic in the recent work looking at direct updating of parameters in the baseline LVCSR model rather than reranking (Kuo *et al.* 2007; Lin and Yvon 2005). These works adopt an approach to efficient LVCSR that involves building a finite-state ‘decoding graph’ off-line and performing optimizations on this transducer, as described in detail in Mohri *et al.* (2002). The decoding graph is the

Idx	Param	Feature
1	α_1	a
2	α_2	b
3	α_3	<s>a
4	α_4	<s>b
5	α_5	aa
6	α_6	ab
7	α_7	a</s>
8	α_8	ba
9	α_9	bb
10	α_{10}	b</s>

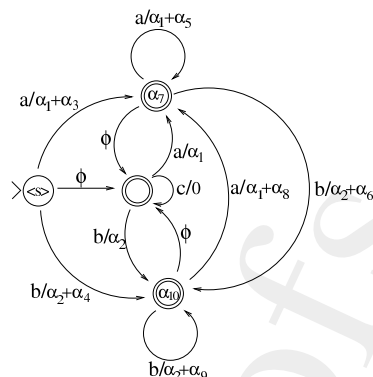


Figure 8.3 Deterministic automaton reading in symbols and assigning parameter weights.

result of combining an automata representation of an n -gram language model (G) with a pronunciation dictionary represented as a finite-state transducer (L) and decision tree models that split the phones in the pronunciations into context-dependent (CD) phone representations (C). The composition of these weighted transducers³ $C \circ L \circ G$ results in a single transducer, which serves to guide decoding. Optimizations to this graph, such as determinization on the input side and weight pushing, can result in a far more efficient search than combining these models on-the-fly during search.

The idea pursued in Lin and Yvon (2005) and extended to full context-dependent LVCSR systems in Kuo *et al.* (2007) is to directly update transition weights within the decoding graph. To do this, they determine the state sequence through the decoding graph of the reference transcription by optimally force aligning the transcription with the speech. Training in Kuo *et al.* (2007) is an on-line procedure discussed in section 8.2.3 that compares the best scoring state sequence with the state sequence derived from the reference, and updates the transition weights accordingly. Hence each transition weight in the graph is an independent feature within a log linear framework.

This approach to feature mapping brings up some interesting questions. First, what is discriminative language modeling? In this instance, the features include context-dependent phone identities as part of the state definition, and stochastic modeling of such features is typically outside the range of language modeling. While it is not the intent of this chapter to try to philosophically delimit the scope of this topic, it seems clear that the transition weights in these graphs represent discrete symbol sequences rather than the Gaussian densities modeled at the states, hence are certainly modeling the language not the acoustics. However, it highlights the fact that there is little in this general framework to limit the kinds of features that might be used, and that future work may find hybrid feature sets – e.g. for combining syntactic and prosodic features – that will lead to substantial system improvements.

The second question that arises is, what do these features represent? The input stream of the transducer is labeled with context-dependent phones, and the output stream is labeled with words in the vocabulary, so in its simplest form, these features represent (CD phone \times

³Note that an automaton is a special case of a transducer where every symbol maps to itself.

word) sub-sequences. These graphs, however, are optimized to increase the amount of shared structure between words with similar CD phone realizations, hence the transitions themselves may represent CD phones over sets of n -grams. These complicated composite features do appear to have strong utility for this task (see Kuo *et al.* 2007 for empirical results), and they capture key distinctions that are clearly lost in most language modeling approaches. The downside of this approach is a loss of generalization from one phone sequence realization of a word sequence to another realization of the same word sequence. One could imagine an extension of this approach that treats arcs as features, but also includes raw n -gram features, and updates arcs in the graph with respect to all relevant features. In such a way, the benefits of this approach could be achieved while also gaining generalization to other pronunciations of the same n -gram.

8.2.3 Parameter Estimation

The final component to the general framework is the parameter estimation algorithm, for choosing a parameter weight vector $\alpha \in \mathbb{R}^d$. Here any technique that has been used for log linear modeling in other domains can be used, with the proviso that the method must scale to very large problems. Language modeling for LVCSR is typically most effective with a large number of candidates (at least in the hundreds) and a very large number of features (in the millions). This means that there have been limited uses of resource intensive methods such as Support Vector Machines (SVMs) (though see section 8.3.3 for one example using SVMs).

Before getting into specific training objectives, it is probably worth spending a little time to address a potential confusion of terminology. In language modeling, *conditional* models have typically signified the conditional probabilities used via the chain rule to define the joint probability of a string. For example, a standard trigram model is defined as the probability of a word conditioned on the previous two words. In the current context, however, conditional models do not denote the local conditional probabilities that are used to define the joint generative probability, but rather denote the likelihood of the string *given* the input utterance. This sort of conditional model is discriminative.

Given the popularity of Maximum Mutual Information Estimation (MMIE) for acoustic modeling (Bahl *et al.* 1986), which has a conditional likelihood objective in the sense we are using here, direct applications of such techniques have been pursued for language modeling. As in prior work of Stolcke and Weintraub (1998), Chueh *et al.* (2005) examines a log linear model which maximizes an unregularized conditional likelihood objective. Linguistic (n -gram) features of the model are combined with a baseline acoustic model score, which can be thought of as just another feature in the model. If \bar{y}_i^{ref} is the reference output for speech input \bar{x}_i , then the conditional log likelihood of the training data of m utterances given a feature mapping ψ and parameter vector α is

$$\begin{aligned}
 \text{LogLike}(\alpha) &= \sum_{i=1}^m \log P(\bar{y}_i^{\text{ref}} | \bar{x}_i) \\
 &= \sum_{i=1}^m \log \left(\frac{\exp(\psi(\bar{x}_i, \bar{y}_i^{\text{ref}}) \cdot \alpha)}{\sum_{\bar{y} \in \text{GEN}(\bar{x}_i)} \exp(\psi(\bar{x}_i, \bar{y}) \cdot \alpha)} \right) \\
 &= \sum_{i=1}^m \left(\psi(\bar{x}_i, \bar{y}_i^{\text{ref}}) \cdot \alpha - \log \sum_{\bar{y} \in \text{GEN}(\bar{x}_i)} \exp(\psi(\bar{x}_i, \bar{y}) \cdot \alpha) \right). \quad (8.2)
 \end{aligned}$$

01 This is a convex function, and hence the α that maximizes the log likelihood can be found
 02 through standard techniques relying on the gradient (see below).

03 One key problem with this approach is overtraining, which, as usual with language
 04 modeling, is severe and has led to little or no system gain for approaches pursuing it as
 05 presented above (Chueh *et al.* 2005; Stolcke and Weintraub 1998). With millions of features,
 06 the opportunity for overtraining is much greater for language modeling than for acoustic
 07 modeling, hence some method for controlling for overtraining (regularization) needs to be
 08 added to the training approach in order to improve generalization to new data. Of course,
 09 the reasons why these models provided little or no gain in system performance may also
 10 be related to some of the issues in training data preparation presented in section 8.2.1, but
 11 certainly lack of regularization also contributed. Similar lack of system improvement was
 12 observed in Paciorek and Rosenfeld (2000) when using Minimum Classification Error (MCE)
 13 training (Juang *et al.* 1997) for language model training, as well as in (Kuo and Chen 2005)
 14 when using Minimum Word Error (MWE) training (Povey and Woodland 2002). Direct
 15 application of popular acoustic model training approaches have thus not been successful,
 16 and a large reason for this is almost certainly lack of appropriate model regularization.

17 Adopting regularization techniques from other sequence modeling work making use of
 18 global conditional likelihood optimization (Johnson *et al.* 1999; Lafferty *et al.* 2001), the
 19 regularized objective in Roark *et al.* (2007, 2004) becomes:

$$20 \text{RegLogLike}(\alpha) = \sum_{i=1}^m \left(\psi(\bar{x}_i, \bar{y}_i^{\text{ref}}) \cdot \alpha - \log \sum_{\bar{y} \in \text{GEN}(\bar{x}_i)} \exp(\psi(\bar{x}_i, \bar{y}) \cdot \alpha) - \frac{\|\alpha\|^2}{2\sigma^2} \right), \quad (8.3)$$

21 which is the common zero-mean Gaussian regularizer that imposes a penalty within the
 22 objective as the magnitude of the parameter values grow. This is also a convex function, thus
 23 guaranteeing convergence to the globally optimal solution. The derivative of this function
 24 with respect to any given parameter α_k is

$$25 \frac{\partial \text{RegLogLike}(\alpha)}{\partial \alpha_k} = \sum_{i=1}^m \left(\psi_k(\bar{x}_i, \bar{y}_i^{\text{ref}}) - \sum_{\bar{y} \in \text{GEN}(\bar{x}_i)} P(\bar{y} | \bar{x}_i) \psi_k(\bar{x}_i, \bar{y}) \right) - \frac{\alpha_k}{\sigma^2}. \quad (8.4)$$

26 The empirical results in Roark *et al.* (2007, 2004) validate that the unregularized global
 27 conditional likelihood objective in Equation (8.2) did not lead to any system improve-
 28 ments, while the regularized objective in Equation (8.3) does provide statistically significant
 29 improvements over the baseline, using an empirically chosen regularization parameter σ .
 30 The gradients of the n -gram features in that work were efficiently extracted from large
 31 word lattices using general finite-state techniques, which enabled the approach to be scaled
 32 up to over 10 million features, using lattices encoding a very large number of candidate
 33 transcriptions, for a training corpus consisting of hundreds of hours of speech.

34 Another technique adopted from the recent machine learning literature is the averaged
 35 Perceptron (Collins 2002; Freund and Schapire 1999), which was shown to provide signif-
 36 icant word-error rate reductions over a baseline LVCSR system in Roark *et al.* (2004) and
 37 was later compared with global conditional likelihood optimization in Roark *et al.* (2007,
 38 2004). The reader is referred to the cited papers for a formal presentation of the algorithm.
 39 Informally, this is an on-line algorithm that initializes parameters at zero and updates the
 40 parameters after every example as follows: using the current model, find the best scoring
 41
 42
 43
 44
 45
 46
 47

01 candidate of the current example (\bar{y}_i^{\max}); then update the parameters by penalizing features
 02 extracted from \bar{y}_i^{\max} and rewarding features extracted from \bar{y}_i^{ref} , the gold standard. Thus, the
 03 parameter update at example $(\bar{x}_i, \bar{y}_i^{\text{ref}})$, in its simplest form, is

$$05 \quad \alpha[i] = \alpha[i - 1] + \psi(\bar{x}_i, \bar{y}_i^{\text{ref}}) - \psi\left(\bar{x}_i, \arg \max_{\bar{y} \in \text{GEN}(\bar{x}_i)} \psi(\bar{x}_i, \bar{y}) \cdot \alpha[i - 1]\right), \quad (8.5)$$

06 where $\alpha[i - 1]$ is the parameter vector from the previous time step⁴. Of course, if $\bar{y}_i^{\text{ref}} =$
 07 $\arg \max_{\bar{y} \in \text{GEN}(\bar{x}_i)} \psi(\bar{x}_i, \bar{y}) \cdot \alpha[i - 1]$, then the parameters are unchanged.

08 Like the conditional likelihood optimization, this approach is prone to overtraining,
 09 hence requires some regularization. The most common technique used for these models is
 10 *averaging* the parameters over the models at every timestep of the algorithm. The averaged
 11 parameter $\text{avg}_i(\alpha_k)$ at time i is defined as

$$12 \quad \text{avg}_i(\alpha_k) = \frac{1}{i} \sum_{j=1}^i \alpha_k[j]. \quad (8.6)$$

13 The raw parameter values are used while training the model, but this averaged parameter
 14 value is then used for evaluating the model on held aside or test data. Typically, some held-
 15 aside data is evaluated after every iteration over the training data, to determine when to stop
 16 training.

17 One benefit of Perceptron training versus regularized conditional likelihood optimization
 18 is that the Perceptron training typically arrives at a more parsimonious solution, i.e. the
 19 number of features with non-zero parameter weights are fewer, at least with commonly used
 20 regularizers. The reason for this is clear – parameters can only be updated in the Perceptron
 21 algorithm if the feature is extracted from either the gold standard candidate \bar{y}_i^{ref} or the best
 22 scoring candidate \bar{y}_i^{\max} , hence many features never move away from their initial values of
 23 zero. In the comparisons between Perceptron and conditional likelihood optimization in
 24 Roark *et al.* (2007, 2004), the conditional likelihood estimation resulted in larger gains over
 25 the baseline than the Perceptron estimation, but the Perceptron was useful as an initial step
 26 prior to the conditional likelihood optimization, both for initializing parameter weights for
 27 fast convergence and for selecting features to include in the model.

28 Recent work has looked at improving on Perceptron language model training by intro-
 29 ducing loss-sensitivity to the on-line parameter update in Equation (8.5). In the approach
 30 of Singh-Miller and Collins (2007), the contribution of an example to the change in the
 31 parameter values is based on a loss value, which in the cited paper is taken as the number of
 32 errors beyond the minimum error rate hypotheses in the n -best lists. Features extracted from
 33 highly errorful hypotheses will be more heavily penalized. In addition, this approach provides
 34 a well-motivated way of dealing with multiple candidates with the same minimum error rate
 35 in the set of candidates. The previously discussed (section 8.2.2) work in Kuo *et al.* (2007)
 36 also uses an on-line parameter estimation approach related to the Perceptron algorithm, where
 37 the magnitude of the parameter value update is determined via a sigmoid function.

38 While the above-mentioned parameter estimation techniques make use of well behaved
 39 (convex) functions, the objectives are not precisely what is typically evaluated, which is

40 ⁴The algorithm may perform multiple passes over the training data, so the index i may indicate the number of
 41 updates, rather than the absolute example number in the corpus.

01 word-error rate (WER): the number of insertions, deletions or substitutions per 100 words
02 of reference text. The global conditional likelihood approach is maximizing the (regularized)
03 conditional likelihood of the gold standard; and the Perceptron is optimizing the string
04 error rate rather than the WER. MWE training (Povey and Woodland 2002) optimizes a
05 continuous function more closely related to error rate than either of these two objective
06 functions, so perhaps a regularized version of that would be worth investigating. Others
07 have tried methods of optimizing the WER function directly, which cannot be done with
08 any guarantees of convergence or finding a global optimum. For example, minimum sample
09 risk training (Gao *et al.* 2005) performs a greedy search of parameterizations by performing
10 a line search along each dimension in sequence. Although no guarantees are given for finding
11 an optimal solution, reported empirical results (Gao *et al.* 2005) are competitive with other
12 methods such as Boosting and the averaged Perceptron. Similarly unguaranteed parameter
13 search techniques were used in Banerjee *et al.* (2003) to update bigram parameter weights
14 for word-error reduction.

15 We have presented a basic, general characterization of how these sorts of problems
16 can be approached, and have pointed to specific papers pursuing specific versions of this
17 general approach. Some of these have achieved relatively large improvements (more than 1%
18 absolute) over baseline LVCSR systems, using just n -gram features. In the following section,
19 we will survey some extensions to the basic approach, in three key directions: novel features;
20 novel objective functions; and within the context of domain adaptation.

21

22

23 **8.3 Further Developments**

24

25 **8.3.1 Novel Features**

26

27 Up to now, we have discussed models that rely on n -grams as features alongside the baseline
28 model score, or treat arcs within a decoding graph as features, e.g. Kuo *et al.* (2007). In
29 this section, we point to three relatively recent papers that look to expand upon these simple
30 feature sets to achieve additional discrimination. Note that this sort of enterprise has been
31 notoriously difficult in generative language modeling for LVCSR, where WER reductions
32 due to features other than n -grams have been hard to come by.

33 Syntactic language models have been a topic of some interest over the past decade, with
34 some notable successes in achieving WER reductions, e.g. Charniak (2001); Chelba and
35 Jelinek (2000); Roark (2001); Wang *et al.* (2003). In these works, generative context-free
36 (Charniak 2001; Chelba and Jelinek 2000; Roark 2001) or finite-state (Wang *et al.* 2003)
37 models are used to define a distribution over strings that can be used as a language model.
38 Within a discriminative language modeling framework, however, Collins *et al.* (2005) used
39 a generative context-free parsing model not to derive a score, but rather to derive a syntactic
40 annotation which was then used within the feature mapping ψ for discriminative modeling
41 with the Perceptron algorithm. Example features included part-of-speech (POS) and shallow
42 parse sub-sequence features; context-free rule instances in the parses; and bilexical head-
43 to-head dependencies. This stands in marked contrast to the generative models, where the
44 generative model scores were used for modeling and the structural annotations were an
45 unused byproduct. The WER reductions in Collins *et al.* (2005) were statistically significant
46 (though modest) versus just using n -gram features, and most of the reduction was achieved
47 with POS-tag sequence features.

01 A similar approach was taken for incorporating morphological features into language
02 models in Czech (Shafran and Hall 2006). Unlike generative morphology-based language
03 models, which have been used profitably in, for example, Arabic speech recognition
04 (Vergyri *et al.* 2004), the use of the kinds of sub-word features that morphological analysis
05 affords does not necessitate complicated smoothing techniques. Rather, the annotations from
06 morphological analyzers can be used within the feature mapping to provide additional
07 features in the model. This is particularly important in highly inflected or agglutinative
08 languages, and promises substantial improvements in such languages, as demonstrated in
09 Shafran and Hall (2006).

10 One key issue in discriminative syntactic and morphological language modeling
11 approaches is feature selection. Features can be derived from rich structural annotations in
12 many ways, which can lead to an astronomical number of potential features in the model.
13 Features were selected in Shafran and Hall (2006) via a χ^2 statistic derived from the co-
14 occurrence of features and the classes of interest, in this case ‘correct’ and ‘incorrect’
15 candidates. A large χ^2 statistic indicates a high correlation between a feature and either the
16 set of correct candidates or the set of incorrect candidates, hence indicating that the feature
17 would have high discriminative utility.

18 Trigger features (Lau *et al.* 1993) are a way of capturing the bursty and topical nature
19 of language – if a word is observed, then there is typically a high likelihood of that word
20 occurring again or related words occurring. A self-trigger is a feature that exploits an
21 instance of a word or term to increase the likelihood of observing another, or in the case
22 of discriminative models, to reward candidates that have the triggered term. In Singh-Miller
23 and Collins (2007), self-triggers of various sorts were used within a loss-sensitive Perceptron
24 modeling framework to reduce WER significantly over using just n -gram features. To do
25 so, they define a history h of previous utterances in the conversation, and derive trigger
26 features based on unigrams and bigrams that either a) occur multiple times in the candidate
27 transcription; or b) occur in the candidate transcription and in transcriptions in the history h .
28 In addition, they create backoff bins of unigrams and bigrams based on TF-IDF statistics, and
29 derive another set of features based on these bins, as a way of tying the parameters for words
30 that are similarly distributed. The combination of raw and binned trigger features, along with
31 standard n -gram features, resulted in the best performing system.

32 Real gains are being achieved by expanding the feature set beyond simple n -gram features,
33 and the discriminative language modeling paradigm that we have been presenting is a natural
34 framework for work on novel methods for feature extraction and selection.

35

36

37 **8.3.2 Novel Objectives**

38 While most research in LVCSR over the years has focused on WER as an objective,
39 it is increasingly becoming the case that speech recognition is a sub-process within a
40 larger application such as call-routing, spoken language understanding or spoken document
41 retrieval. Within such a scenario, discriminative language modeling approaches can be
42 applied with a downstream objective function in training rather than an objective function
43 based on the accuracy of the speech recognition itself. This sort of optimization can occur
44 within the context of generative language model estimation, as with the spoken language
45 understanding applications in Goel (2004); Goel *et al.* (2005), in which bigram probabilities
46 were iteratively re-estimated to improve utterance annotation (Goel 2004) or call-routing
47

01 (Goel *et al.* 2005) performance. Discriminative language modeling methods, however, allow
02 for great flexibility of feature mapping along with direct optimization of such an objective.

03 In Saraclar and Roark (2005, 2006), utterances within a call-routing application were
04 annotated with class labels (97 classes), and discriminative language model estimation
05 was done using either the Perceptron algorithm (Saraclar and Roark 2005) or regularized
06 global conditional likelihood maximization (Saraclar and Roark 2006) for various objective
07 functions: a WER objective; an utterance class error rate (CER) objective; and a joint
08 WER/CER objective. Features included n -grams, class labels, and combination class label
09 and n -gram. Optimization of the joint objective led to a very slight (not statistically
10 significant) increase of CER versus training with an exclusively CER objective, but no
11 increase in WER was observed versus training with an exclusively WER objective. That is,
12 the trade-off inherent in a joint optimization came at the expense of CER rather than WER,
13 which is unsurprising given that each utterance consists of many words but receives just one
14 class label. Inclusion of the utterance class in the model achieved modest (though statistically
15 significant) improvements in WER versus just n -gram features, hence this can also be seen
16 as a method for increasing the feature set.

17 Given the growing interest in machine translation (MT) of speech, one might expect future
18 work in pushing MT objective functions back into LVCSR model training, and discriminative
19 language modeling of the sort described here is one way to achieve this. For example, the
20 presence or absence of phrases from a phrase table within a phrase-based MT system may
21 impact the final quality of translation, hence optimizing the LVCSR system to do a better job
22 of recognizing such phrases, perhaps at the expense of overall WER, which may result in an
23 overall system improvement. Similar approaches could be pushed in speech IR or other areas Q64
24 of spoken language understanding. Other than the papers discussed here, we are unaware of
25 any published results in this direction.

26

27

28

8.3.3 Domain Adaptation

29

30 All of the work discussed up to now has been within the context of in-domain trained systems,
31 but language model adaptation to novel domains is an area of great interest where they have
32 also been shown to have utility. An adaptation scenario does change a couple of key aspects
33 of the training that impact the best practices. First, to the extent that the new domain is not
34 part of the baseline training data (which is the whole point of domain adaptation) then the
35 issue of cross validating the training set to get training data is moot. However, there remains
36 the issue of the reference transcription not being in the lattice output of the recognizer.
37 In fact, the oracle accuracy of the lattice, i.e. the best possible accuracy achievable from
38 among the candidates, will likely be low for the novel domain, which can reduce the efficacy
39 of these discriminative approaches, relative to generative alternatives such as Maximum a
40 Posteriori (MAP) adaptation (Bacchiani *et al.* 2006; Gauvain and Lee 1994).

41 In Bacchiani *et al.* (2004), a large vocabulary general voicemail transcription system was
42 adapted for customer service support voicemail recognition. The baseline out-of-domain
43 training of roughly 100 hours was augmented with 17 hours of in-domain training data.
44 In a straight head-to-head comparison, MAP adaptation of the generative language model,
45 based on counts derived from the in-domain data, led to significantly larger improvements
46 in transcription accuracy over the baseline (nearly 8% absolute) than using the Perceptron
47 algorithm (5% absolute). Note that the vocabulary was kept constant in this application, so the

01 MAP adaptation did not benefit from additional in-domain vocabulary. The hypothesis was
02 that the re-ranking nature of the Perceptron was hurting it, since the MAP adapted generative
03 model was used in the first pass, hence could find good candidate transcriptions that were not
04 in the word lattices produced by the baseline out-of-domain trained recognizer.

05 To remedy this, a hybrid of MAP adaptation followed by Perceptron reranking was
06 pursued in Bacchiani *et al.* (2004). Of course, once the in-domain data was included in
07 the baseline LVCSR training data, appropriate cross-validation techniques had to be used to
08 produce training lattices for the discriminative language model. The result of this approach
09 was an additional 0.7% absolute improvement versus using MAP adaptation alone.

10 In Zhou *et al.* (2006), several discriminative training algorithms were compared for
11 domain adaptation in Chinese LVCSR, including the Perceptron algorithm, boosting
12 (Schapire *et al.* 1998), ranking SVMs (Joachims 2002) and minimum sample risk (Gao *et al.*
13 2005, see section 8.2.3). They found the Perceptron algorithm provided the best domain
14 adaptation (by more than 1% absolute), which they attributed to the large size of the feature
15 set chosen by the Perceptron algorithm relative to the other methods investigated. In addition,
16 they found that the ranking SVM took two orders of magnitude more training time than the
17 other methods.

18 What remain unexplored are semi-supervised adaptation techniques for this problem –
19 making use of in-domain text without recordings or recordings without transcriptions to
20 adapt language models using discriminative parameter estimation techniques. The flexibility
21 of generative adaptation approaches, such as MAP estimation, permits the opportunistic
22 exploitation of available resources for model training. Achieving such flexibility within
23 discriminative language modeling approaches is an important topic that deserves further
24 attention.

25
26

27 8.4 Summary and Discussion

28
29 In this chapter, we have presented a general discriminative modeling framework and its appli-
30 cation to language modeling for LVCSR. Nearly all of the work on discriminative language
31 modeling fits this general characterization, and it is useful for defining the dimensions within
32 which approaches differ. We have made an effort to be comprehensive in the survey, within
33 the scope declared at the outset of the chapter. Overall the number of papers that have been
34 written on this topic are relatively few, which can be attributed to a number of factors.

35 First, there are significant barriers to entry in researching this topic. Unlike other
36 application areas, such as machine translation, achieving state-of-the-art performance in
37 LVCSR is not accessible to researchers outside of a small number of research labs and
38 centers. The number of techniques that must be combined within a full-blown LVCSR
39 system to achieve this performance – including a range of acoustic modeling techniques in
40 multiple stages, first-pass large vocabulary decoding, speaker adaptation and lattice rescoring
41 techniques – are such that a researcher must expect a major time investment to even produce
42 the training data needed to explore discriminative language modeling. What should be a
43 priority within the speech community is to lower these barriers by making a large amount
44 of training and evaluation word lattices available for research purposes to the community at
45 large. That would almost certainly spur researchers in machine learning and NLP to begin
46 applying their work in this domain.
47

01 It used to be the case that perplexity reduction was the primary mode for evaluating
02 language model improvements. Given the poor correlation between reductions in perplexity
03 and WER reduction, this is no longer widely accepted as a language modeling objective.
04 It did have the virtue, however, of allowing evaluation on text corpora, which meant that
05 research on language modeling could be conducted without access to LVCSR systems. For
06 the discriminative models presented here, this is no longer the case: LVCSR systems are
07 required for training data, and WER is the objective – perplexity cannot be calculated from
08 these models. One step towards increasing research on language modeling would be to make
09 word lattices from competitive LVCSR systems available to the community at large.

10 Even if such data were made available, however, this remains a very large and difficult
11 problem. Scaling up machine learning algorithms to handle problem spaces with many
12 millions of dimensions and training sets with millions or billions of examples (or more)
13 is a challenge. There is a reason for discussing finite-state methods in this survey, and that
14 is the scale of the problem. Coupled with the widespread perception that language modeling
15 for LVCSR is hopeless, these challenges present another kind of barrier to entry.

16 This is a shame, because there has been progress in language modeling in this direction,
17 and there remain a large number of very interesting and relatively unexplored directions in
18 which to take it, such as the topics in the following short list:

- 19 • Semi-supervised techniques for leveraging large text corpora alongside the kind of
20 training data we have been discussing;
- 21
- 22 • Further research on richer feature sets, such as syntactic features or hybrid acous-
23 tic/language features;
- 24
- 25 • Feature selection and/or dimensionality reduction;
- 26
- 27 • Scaling up compute intensive parameter estimation techniques.

28 This survey suggests that it is not unreasonable to foster the hope that real LVCSR system
29 improvements can and will be achieved along these lines over the next decade, to the point
30 where resulting techniques will become part of the collective best practices in the field.

31

32 References

33

- 34 Allauzen C, Mohri M and Roark B 2003 Generalized algorithms for constructing language models.
35 *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics (ACL)*,
36 pp. 40–47.
- 37 Bacchiani M, Riley M, Roark B and Sproat R 2006 MAP adaptation of stochastic grammars. *Computer
38 Speech and Language* **20**(1), 41–68.
- 39 Bacchiani M, Roark B and Saraclar M 2004 Language model adaptation with MAP estimation and
40 the perceptron algorithm. *Proceedings of the Human Language Technology Conference of the North
41 American Chapter of the Association for Computational Linguistics (HLT-NAACL)*, pp. 21–24.
- 42 Bahl L, Brown P, de Souza P and Mercer R 1986 Maximum mutual information estimation of hidden
43 Markov model parameters for speech recognition. *Proceedings of the IEEE International Conference
44 on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 49–52.
- 45 Bahl L, Brown P, de Souza P and Mercer R 1993 Estimating hidden Markov model parameters so as
46 to maximize speech recognition accuracy. *IEEE Transactions on Speech and Audio Processing* **1**(1),
47 77–83.

Q65

- 01 Banerjee S, Mostow J, Beck J and Tam W 2003 Improving language models by learning from speech
02 recognition errors in a reading tutor that listens. *Proceedings of the 2nd International Conference on*
03 *Applied Artificial Intelligence*, Fort Panhala, Kolhapur, India.
- 04 Charniak E 2001 Immediate-head parsing for language models. *Proceedings of the 39th Annual*
05 *Meeting of the Association for Computational Linguistics (ACL)*.
- 06 Chelba C and Jelinek F 2000 Structured language modeling. *Computer Speech and Language* **14**(4),
07 283–332.
- 08 Chen Z, Lee KF and Li MJ 2000 Discriminative training on language model. *Proceedings of the*
09 *International Conference on Spoken Language Processing (ICSLP)*.
- 10 Chueh C, Chien T and Chien J 2005 Discriminative maximum entropy language model for speech
11 recognition. *Proceedings of the European Conference on Speech Communication and Technology*
12 *(Interspeech)*.
- 13 Collins M 2002 Discriminative training methods for hidden Markov models: Theory and experiments
14 with perceptron algorithms. *Proceedings of the Conference on Empirical Methods in Natural*
15 *Language Processing (EMNLP)*, pp. 1–8.
- 16 Collins M, Saraclar M and Roark B 2005 Discriminative syntactic language modeling for speech
17 recognition. *Proceedings of the 43rd Annual Meeting of the Association for Computational*
18 *Linguistics (ACL)*, pp. 507–514.
- 19 Collins MJ 2000 Discriminative reranking for natural language parsing. *Proceedings of the 17th*
20 *International Conference on Machine Learning*.
- 21 Emami A, Papineni K and Sorensen J 2007 Large-scale distributed language modeling. *Proceedings of*
22 *the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- 23 Freund Y and Schapire R 1999 Large margin classification using the perceptron algorithm. *Machine*
24 *Learning* **3**(37), 277–296.
- 25 Gao J, Yu H, Yuan W and Xu P 2005 Minimum sample risk methods for language modeling.
26 *Proceedings of the Conference on Human Language Technology Conference and Empirical Methods*
27 *in Natural Language Processing (HLT-EMNLP)*, pp. 209–216.
- 28 Gauvain JL and Lee CH 1994 Maximum a posteriori estimation for multivariate Gaussian mixture
29 observations of Markov chains. *IEEE Transactions on Speech and Audio Processing* **2**(2), 291–298.
- 30 Goel V 2004 Conditional maximum likelihood estimation for improving annotation performance of
31 n-gram models incorporating stochastic finite state grammars. *Proceedings of the International*
32 *Conference on Spoken Language Processing (ICSLP)*.
- 33 Goel V and Byrne W 2000 Minimum Bayes-risk automatic speech recognition. *Computer Speech and*
34 *Language* **14**(2), 115–135.
- 35 Goel V, Kuo H, Deligne S and Wu C 2005 Language model estimation for optimizing end-to-end
36 performance of a natural language call routing system. *Proceedings of the IEEE International*
37 *Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. I/565–568.
- 38 Goodman J 2001 A bit of progress in language modeling. *Computer Speech and Language* **15**(4),
39 403–434.
- 40 Jelinek F 1996 Acoustic sensitive language modeling. Center for Language and Speech Processing,
41 Johns Hopkins University, Baltimore, MD. Technical report.
- 42 Joachims T 2002 Optimizing search engines using clickthrough data. *Proceedings of the ACM*
43 *Conference on Knowledge Discovery and Data Mining (KDD)*.
- 44 Johnson M, Geman S, Canon S, Chi Z and Riezler S 1999 Estimators for stochastic “unification-based”
45 grammars. *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*
46 *(ACL)*, pp. 535–541.
- 47 Juang B, Chou W and Lee C 1997 Minimum classification error rate methods for speech recognition.
IEEE Transactions on Speech and Audio Processing **5**(3), 257–265.

- 01 Khudanpur S and Wu J 2000 Maximum entropy techniques for exploiting syntactic, semantic and
02 collocational dependencies in language modeling. *Computer Speech and Language* **14**(4), 355–372.
- 03 Kuo H, Kingsbury B and Zweig G 2007 Discriminative training of decoding graphs for large vocabulary
04 continuous speech recognition. *Proceedings of the IEEE International Conference on Acoustics,
05 Speech, and Signal Processing (ICASSP)*, pp. IV/45–48.
- 06 Kuo HKJ, Fosler-Lussier E, Jiang H and Lee CH 2002 Discriminative training of language models for
07 speech recognition. *Proceedings of the IEEE International Conference on Acoustics, Speech, and
08 Signal Processing (ICASSP)*, Orlando, FL.
- 09 Kuo J and Chen B 2005 Minimum word error based discriminative training of language models.
10 *Proceedings of the European Conference on Speech Communication and Technology (Interspeech)*.
- 11 Lafferty J, McCallum A and Pereira F 2001 Conditional random fields: Probabilistic models for
12 segmenting and labeling sequence data. *Proceedings of the 18th International Conference on
13 Machine Learning*, pp. 282–289.
- 14 Lau R, Rosenfeld R and Roukos S 1993 Trigger-based language models: a maximum entropy approach.
15 *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing
16 (ICASSP)*, pp. 45–48.
- 17 Lin S and Yvon F 2005 Discriminative training of finite-state decoding graphs. *Proceedings of the
18 European Conference on Speech Communication and Technology (Interspeech)*.
- 19 Mangu L and Padmanabhan M 2001 Error corrective mechanisms for speech recognition. *Proceedings
20 of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- 21 Mangu L, Brill E and Stolcke A 2000 Finding consensus in speech recognition: word error minimization
22 and other application of confusion networks. *Computer Speech and Language* **14**(4), 373–400.
- 23 Mohri M and Riley M 1997 Weighted determinization and minimization for large vocabulary speech
24 recognition. *Proceedings of the European Conference on Speech Communication and Technology
25 (Interspeech)*.
- 26 Mohri M, Pereira FCN and Riley M 2002 Weighted finite-state transducers in speech recognition.
27 *Computer Speech and Language* **16**(1), 69–88.
- 28 Paciorek C and Rosenfeld R 2000 Minimum classification error training in exponential language
29 models. *Proceedings of the NIST Speech Transcription Workshop*.
- 30 Popat TBA, Xu P, Och F and Dean J 2007 Large language models in machine translation. *Proceedings
31 of the Joint Conference on Empirical Methods in Natural Language Processing and Computational
32 Natural Language Learning (EMNLP-CoNLL)*, pp. 858–867.
- 33 Povey D and Woodland P 2002 Minimum phone error and I-smoothing for improved discriminative
34 training. *Proceedings of the International Conference on Spoken Language Processing (ICSLP)*.
- 35 Ringger EK and Allen JF 1996 Error corrections via a post-processor for continuous speech recognition.
36 *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing
37 (ICASSP)*.
- 38 Roark B 2001 Probabilistic top-down parsing and language modeling. *Computational Linguistics* **27**(2),
39 249–276.
- 40 Roark B, Saraclar M and Collins M 2004 Corrective language modeling for large vocabulary ASR with
41 the perceptron algorithm. *Proceedings of the IEEE International Conference on Acoustics, Speech,
42 and Signal Processing (ICASSP)*, pp. I/749–752.
- 43 Roark B, Saraclar M and Collins M 2007 Discriminative n-gram language modeling. *Computer Speech
44 and Language* **21**(2), 373–392.
- 45 Roark B, Saraclar M, Collins M and Johnson M 2004 Discriminative language modeling with
46 conditional random fields and the perceptron algorithm. *Proceedings of the 42nd Annual Meeting
47 of the Association for Computational Linguistics (ACL)*, pp. 47–54.
- Rosenfeld R, Chen S and Zhu X 2001 Whole-sentence exponential language models: a vehicle for
linguistic-statistical integration. *Computer Speech and Language* **15**(1), 55–73.

- 01 Saraclar M and Roark B 2005 Joint discriminative language modeling and utterance classification.
02 *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*
03 *(ICASSP)*, pp. I/561–564.
- 04 Saraclar M and Roark B 2006 Utterance classification with discriminative language modeling. *Speech*
05 *Communication* **48**(3-4), 276–287.
- 06 Schapire RE, Freund Y, Bartlett P and Lee WS 1998 Boosting the margin: A new explanation for the
07 effectiveness of voting methods. *The Annals of Statistics* **26**(5), 1651–1686.
- 08 Shafran I and Byrne W 2004 Task-specific minimum Bayes-risk decoding using learned edit distance.
09 *Proceedings of the International Conference on Spoken Language Processing (ICSLP)*.
- 10 Shafran I and Hall K 2006 Corrective models for speech recognition of inflected languages. *Proceedings*
11 *of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 390–398.
- 12 Singh-Miller N and Collins M 2007 Trigger-based language modeling using a loss-sensitive perceptron
13 algorithm. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal*
14 *Processing (ICASSP)*, pp. IV/25–28.
- 15 Stolcke A and Weintraub M 1998 Discriminative language modeling. *Proceedings of the 9th Hub-5*
16 *Conversational Speech Recognition Workshop*.
- 17 Stolcke A, Bratt H, Butzberger J, Franco H, Gadde VRR, Plauche M, Richey C, Shriberg E, Sonmez K,
18 Weng F and Zheng J 2000 The SRI March 2000 Hub-5 conversational speech transcription system.
19 *Proceedings of the NIST Speech Transcription Workshop*.
- 20 Vergyri D, Kirchoff K, Duh K and Stolcke A 2004 Morphology-based language modeling for Arabic
21 speech recognition. *Proceedings of the International Conference on Spoken Language Processing*
22 *(ICSLP)*.
- 23 Wang W, Harper MP and Stolcke A 2003 The robustness of an almost-parsing language model given
24 errorful training data. *Proceedings of the IEEE International Conference on Acoustics, Speech, and*
25 *Signal Processing (ICASSP)*.
- 26 Zhou Z, Gao J, Soong F and Meng H 2006 A comparative study of discriminative methods for
27 reranking LVCSR n-best hypotheses in domain adaptation and generalization. *Proceedings of the*
28 *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. I/141–
29 144.
- 30 Zhou Z, Meng H and Lo W 2006 A multi-pass error detection and correction framework for Mandarin
31 LVCSR. *Proceedings of the International Conference on Spoken Language Processing (ICSLP)*.
- 32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47